



7210 Service Access System

Release 26.9.R1

7210 SAS-Mxp, R6, R12, S, Sx, T MPLS Guide

3HE 27836 AAAB TQZZA 01

Edition: 01

September 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.....	11
List of figures.....	14
1 Getting started.....	16
1.1 About this guide.....	16
1.1.1 Document structure and content.....	17
1.2 7210 SAS modes of operation.....	17
1.3 7210 SAS port modes.....	19
1.4 7210 SAS router configuration process.....	21
1.5 Conventions.....	22
1.5.1 Precautionary and information messages.....	22
1.5.2 Options or substeps in procedures and sequential workflows.....	22
2 MPLS and RSVP.....	24
2.1 MPLS.....	24
2.1.1 MPLS label stack.....	24
2.1.1.1 Label values.....	25
2.1.2 Label Switching Routers.....	26
2.1.2.1 LSP types.....	26
2.1.2.2 MPLS Fast Re-Route (FRR).....	27
2.1.2.3 Manual bypass LSP.....	28
2.1.3 Configuration guidelines.....	31
2.2 MPLS pseudowire hash label support.....	31
2.3 MPLS Transport Profile (MPLS-TP).....	32
2.3.1 MPLS-TP model.....	33
2.3.2 MPLS-TP provider edge and gateway.....	33
2.3.2.1 VLL services.....	33
2.3.3 Detailed descriptions of MPLS-TP.....	35
2.3.3.1 MPLS-TP LSPs.....	35
2.3.3.2 MPLS-TP on pseudowires.....	36
2.3.4 MPLS-TP maintenance identifiers.....	36
2.3.4.1 Generic Associated Channel.....	39
2.3.4.2 MPLS-TP Operations, Administration and Maintenance (OAM).....	40

2.3.4.3	PW control channel status notifications (static pseudowire status signaling).....	43
2.3.4.4	Pseudowire redundancy and active / standby dual-homing.....	44
2.3.4.5	MPLS-TP LSP protection.....	45
2.3.5	Configuring MPLS-TP.....	47
2.3.5.1	Configuration overview.....	47
2.3.5.2	Node-wide MPLS-TP parameter configuration.....	48
2.3.5.3	Node-wide MPLS-TP identifier configuration.....	49
2.3.5.4	Static LSP and pseudowire (VC) label and tunnel ranges.....	49
2.3.5.5	Interface configuration for MPLS-TP.....	50
2.3.5.6	LER configuration for MPLS-TP.....	51
2.3.5.7	Intermediate LSR configuration for MPLS-TP LSPs.....	56
2.4	RSVP.....	57
2.4.1	Using RSVP for MPLS.....	58
2.4.1.1	RSVP Traffic Engineering extensions for MPLS.....	58
2.4.2	Reservation styles.....	59
2.4.2.1	RSVP message pacing.....	60
2.4.3	RSVP overhead refresh reduction.....	60
2.4.3.1	Configuring implicit null.....	60
2.4.4	Using unnumbered Point-to-Point interface in RSVP.....	61
2.4.4.1	Operation of RSVP FRR facility backup over unnumbered interface.....	62
2.4.5	PCEP support for RSVP-TE LSPs.....	63
2.5	Traffic Engineering.....	63
2.5.1	TE metric (IS-IS and OSPF).....	64
2.5.1.1	Maintenance of TE links and nodes.....	64
2.5.2	Admin-group support on facility bypass backup LSP.....	64
2.5.2.1	Procedures at head-end node.....	64
2.5.2.2	Procedures at PLR node.....	65
2.5.3	Manual and timer resignal of RSVP-TE bypass LSP.....	67
2.5.3.1	RSVP-TE bypass LSP path SRLG information update in manual and timer resignal MBB.....	68
2.5.3.2	RSVP-TE bypass LSP path administrative group information update in manual and timer resignal MBB.....	70
2.6	Advanced MPLS/RSVP features.....	71
2.6.1	Shared Risk Link Groups.....	71
2.6.1.1	Enabling disjoint backup paths.....	72
2.6.1.2	Static configurations of SRLG memberships.....	73
2.6.2	TE graceful shutdown.....	74

2.6.3	Inter-area TE LSP (ERO expansion method).....	75
2.6.3.1	Automatic ABR node selection for inter-area LSP.....	75
2.6.3.2	Inter-area LSP support of OSPF virtual links.....	78
2.6.3.3	Area border node FRR protection for inter-area LSP.....	78
2.7	Point-to-Multipoint (P2MP) LSP.....	79
2.7.1	Application in video broadcast.....	79
2.7.2	P2MP LSP data plane.....	80
2.7.2.1	Procedures at ingress LER node.....	80
2.7.2.2	Procedures at LSR node.....	81
2.7.2.3	Procedures at branch LSR node.....	81
2.7.2.4	Procedures at egress LER node.....	81
2.7.2.5	Procedures at BUD LSR node.....	81
2.7.3	RSVP control plane in a P2MP LSP.....	82
2.7.4	Forwarding multicast packets over RSVP P2MP LSP in the base router.....	84
2.7.4.1	Procedures at ingress LER node.....	84
2.7.4.2	Procedures with a primary tunnel at egress LER node.....	85
2.7.5	Configuration guidelines for RSVP P2MP LSPs.....	86
2.8	MPLS/RSVP configuration process overview.....	86
2.9	Configuration notes.....	87
2.10	Configuring MPLS and RSVP with CLI.....	87
2.11	MPLS configuration overview.....	87
2.11.1	LSPs.....	87
2.11.2	Paths.....	87
2.11.3	Router interface.....	87
2.11.4	Choosing the signaling protocol.....	88
2.12	Basic MPLS configuration.....	88
2.13	Common configuration tasks.....	89
2.13.1	Configuring global MPLS parameters.....	90
2.13.2	Configuring an MPLS interface.....	90
2.13.3	Configuring MPLS paths.....	91
2.13.4	Configuring an MPLS LSP.....	91
2.13.4.1	Configuring a static LSP.....	92
2.13.5	Configuring manual bypass tunnels.....	92
2.14	Configuring RSVP parameters.....	94
2.14.1	Configuring RSVP message pacing parameters.....	94
2.14.2	Configuring graceful shutdown.....	95

2.15	MPLS configuration management tasks.....	95
2.15.1	Modifying MPLS parameters.....	95
2.15.2	Modifying an MPLS LSP.....	95
2.15.3	Modifying MPLS path parameters.....	96
2.15.4	Modifying MPLS static LSP parameters.....	96
2.15.5	Deleting an MPLS interface.....	97
2.16	RSVP configuration management tasks.....	97
2.16.1	Modifying RSVP parameters.....	97
2.16.2	Modifying RSVP message pacing parameters.....	98
2.16.3	Deleting an interface from RSVP.....	98
2.17	MPLS/RSVP command reference.....	98
2.17.1	Command hierarchies.....	98
2.17.1.1	MPLS commands.....	99
2.17.1.2	MPLS-TP commands.....	100
2.17.1.3	MPLS LSP commands.....	100
2.17.1.4	MPLS-TP LSP commands.....	102
2.17.1.5	MPLS path commands.....	102
2.17.1.6	RSVP commands.....	103
2.17.1.7	Show commands.....	103
2.17.1.8	Tools commands.....	104
2.17.1.9	Clear commands.....	104
2.17.1.10	Debug commands.....	105
2.17.2	Command descriptions.....	106
2.17.2.1	MPLS configuration commands.....	106
2.17.2.2	RSVP configuration commands.....	180
2.17.2.3	Show commands.....	196
2.17.2.4	Tools commands.....	231
2.17.2.5	Clear commands.....	238
2.17.2.6	Debug commands.....	240
3	Label Distribution Protocol.....	255
3.1	Label Distribution Protocol.....	255
3.1.1	LDP and MPLS.....	255
3.1.2	LDP architecture.....	256
3.1.3	Subsystem interrelationships.....	256
3.1.3.1	Memory manager and LDP.....	257

3.1.3.2	Label manager.....	257
3.1.3.3	LDP configuration.....	258
3.1.3.4	Logger.....	258
3.1.3.5	Service manager.....	258
3.1.4	Execution flow.....	258
3.1.4.1	Initialization.....	258
3.1.4.2	Session lifetime.....	258
3.1.5	Label exchange.....	259
3.1.5.1	Other reasons for label actions.....	259
3.1.5.2	Cleanup.....	259
3.1.5.3	Configuring implicit null label.....	259
3.1.5.4	LDP filters.....	260
3.1.6	ECMP support for LDP.....	260
3.1.6.1	Label operations.....	261
3.1.6.2	LDP LSR ECMP hashing.....	261
3.1.7	Link LDP.....	262
3.1.7.1	Targeted LDP.....	263
3.1.8	Unnumbered interface support in LDP.....	263
3.1.8.1	Feature configuration.....	263
3.1.8.2	Operation of LDP over an unnumbered IP interface.....	264
3.1.9	LDP over RSVP tunnels.....	265
3.1.9.1	Signaling and operation.....	267
3.1.9.2	Rerouting around failures.....	268
3.1.10	T-LDP session tracking using BFD.....	269
3.1.10.1	LDP Downstream-on-Demand (DoD).....	270
3.1.11	LDP over RSVP and ECMP.....	270
3.1.12	LDP Fast-Reroute for IS-IS and OSPF prefixes.....	271
3.1.12.1	LDP FRR configuration.....	271
3.1.12.2	LDP FRR procedures.....	272
3.1.13	LDP P2MP support.....	274
3.1.13.1	LDP P2MP configuration.....	274
3.1.13.2	LDP P2MP protocol.....	274
3.1.13.3	Configuration guidelines for P2MP LSPs.....	274
3.1.14	IS-IS and OSPF support for Loop-Free Alternate calculation.....	275
3.1.14.1	Loop-Free Alternate calculation for inter-area/inter-level prefixes.....	277
3.1.14.2	Loop-Free Alternate Shortest Path First (LFA SPF) policies.....	277

3.1.15	Multi-area and multi-instance extensions to LDP.....	278
3.2	LDP IPv6 control and data planes.....	278
3.2.1	LDP Operation in an IPv6 Network.....	278
3.2.2	Link LDP.....	279
3.2.3	Targeted LDP.....	279
3.2.4	FEC resolution.....	280
3.2.5	Resources required to trap LDP control packets to the CPU.....	280
3.2.6	LDP session capabilities.....	281
3.2.7	LDP adjacency capabilities.....	282
3.2.8	Address and FEC distribution.....	283
3.2.9	Controlling IPv6 FEC distribution during an upgrade to 7210 SAS SR OS supporting LDP IPv6.....	285
3.2.10	Handling of duplicate link-local IPv6 addresses in FEC resolution.....	286
3.2.11	IGP and static route synchronization with LDP.....	287
3.2.12	BFD operation.....	287
3.2.13	Services using SDP with an LDP IPv6 FEC.....	288
3.2.14	Mirror services.....	288
3.2.14.1	Configuration at mirror source node.....	289
3.2.14.2	Configuration at mirror destination node.....	289
3.2.15	OAM Support with LDP IPv6.....	289
3.2.15.1	Configuration guidelines for LDP IPv6 OAM tools.....	290
3.2.16	LDP IPv6 Interoperability Considerations.....	290
3.2.16.1	Interoperability with implementations compliant with RFC 7552.....	290
3.2.16.2	Interoperability with implementations compliant with RFC 5036 for IPv4 LDP control plane only.....	291
3.3	LDP process overview.....	291
3.4	Configuring LDP with CLI.....	292
3.5	LDP configuration overview.....	293
3.6	Basic LDP configuration.....	293
3.7	Common configuration tasks.....	293
3.7.1	Enabling LDP.....	293
3.7.2	Configuring FEC originate parameters.....	294
3.7.3	Configuring graceful-restart helper parameters.....	295
3.7.4	Applying export and import policies.....	295
3.7.5	Targeted session parameters.....	296
3.7.6	Interface parameters.....	297
3.7.7	Peer parameters.....	297

3.7.8	LDP signaling and services.....	298
3.8	LDP configuration management tasks.....	298
3.8.1	Disabling LDP.....	299
3.8.2	Modifying targeted session parameters.....	299
3.8.3	Modifying interface parameters.....	299
3.9	LDP command reference.....	300
3.9.1	Command hierarchies.....	300
3.9.1.1	LDP commands.....	300
3.9.1.2	Show commands.....	302
3.9.1.3	Clear commands.....	303
3.9.1.4	Debug commands.....	303
3.9.2	Command descriptions.....	304
3.9.2.1	LDP configuration commands.....	304
3.9.2.2	Show LDP commands.....	345
3.9.2.3	Clear commands.....	389
3.9.2.4	Debug commands.....	391
4	PCEP.....	397
4.1	Introduction to PCEP.....	397
4.2	Base implementation of PCE.....	400
4.3	PCEP session establishment and maintenance.....	402
4.4	PCEP parameters.....	402
4.4.1	PCC configuration.....	403
4.4.2	LSP initiation.....	403
4.4.3	PCC-initiated and PCE-computed or PCE-controlled LSPs.....	404
4.5	PCEP support for RSVP-TE LSPs.....	406
4.5.1	RSVP-TE LSP configuration for a PCC router.....	406
4.5.2	Behavior of the LSP path update.....	407
4.5.2.1	Path update with empty ERO.....	408
4.5.3	Behavior of LSP MBB.....	408
4.5.3.1	PCC-controlled LSPs.....	408
4.5.3.2	PCE-computed LSPs.....	409
4.5.3.3	PCE-controlled LSPs.....	409
4.5.4	Behavior of secondary LSP Paths.....	411
4.5.5	PCE path profile support.....	411
4.6	LSP path diversity and bidirectionality constraints.....	412

4.7	PCEP configuration command reference.....	413
4.7.1	Command hierarchies.....	413
4.7.1.1	PCEP commands.....	414
4.7.1.2	Show commands.....	414
4.7.1.3	Tools commands.....	414
4.7.2	Command descriptions.....	414
4.7.2.1	PCEP commands.....	415
4.7.2.2	Show commands.....	420
4.7.2.3	Tools commands.....	429
5	Standards and protocol support.....	431
5.1	BGP.....	431
5.2	Certificate Management.....	433
5.3	Ethernet.....	433
5.4	EVPN.....	434
5.5	Fast Reroute.....	435
5.6	Internet Protocol (IP) — General.....	435
5.7	IP — Multicast.....	438
5.8	IP — Version 4.....	439
5.9	IP — Version 6.....	440
5.10	IPsec.....	441
5.11	IS-IS.....	442
5.12	Management.....	443
5.13	MPLS — General.....	446
5.14	MPLS — GMPLS.....	446
5.15	MPLS — LDP.....	447
5.16	MPLS — MPLS-TP.....	447
5.17	MPLS — OAM.....	448
5.18	MPLS — RSVP-TE.....	448
5.19	OSPF.....	448
5.20	Pseudowire.....	450
5.21	Quality of Service.....	450
5.22	RIP.....	451
5.23	Timing.....	451
5.24	VPLS.....	452

List of tables

Table 1: Supported modes of operation and configuration methods.....	18
Table 2: Supported port modes by mode of operation.....	20
Table 3: 7210 SAS platforms supporting port modes.....	21
Table 4: Configuration process.....	22
Table 5: Packet/label field description.....	25
Table 6: Mapping from RSVP-TE to MPLS-TP maintenance identifiers.....	38
Table 7: Address types for the downstream mapping TLV.....	41
Table 8: Bypass LSP admin-group constraint behavior.....	65
Table 9: Determination of bypass LSP path optimality.....	69
Table 10: Output fields: If-attribute admin group.....	197
Table 11: Output fields: MPLS bypass-tunnel.....	199
Table 12: Output fields: MPLS interface.....	201
Table 13: Output fields: MPLS label.....	203
Table 14: Output fields: MPLS label-range.....	204
Table 15: Output fields: MPLS LSP.....	207
Table 16: Output fields: MPLS path.....	214
Table 17: Output fields: MPLS SRLG-group.....	216
Table 18: Output fields: MPLS static-LSP.....	218
Table 19: Output fields: MPLS status.....	219
Table 20: Output fields: RSVP interface.....	224
Table 21: Output fields: RSVP neighbor.....	227

Table 22: Output fields: RSVP session.....	229
Table 23: Output fields: RSVP statistics.....	230
Table 24: Output fields: RSVP status.....	230
Table 25: LSR hashing scenarios.....	261
Table 26: Keepalive timeout factor default values.....	316
Table 27: Default values for hello parameters.....	322
Table 28: Output fields: LDP bindings.....	358
Table 29: Output fields: LDP active bindings.....	360
Table 30: Output fields: LDP active IPv4 bindings.....	363
Table 31: Output fields: LDP active IPv6 bindings.....	366
Table 32: Output fields: LDP IPv4 bindings.....	368
Table 33: Output fields: LDP IPv6 bindings.....	370
Table 34: Output fields: LDP discovery.....	372
Table 35: Output fields: LDP interface.....	374
Table 36: Output fields: LDP parameters.....	376
Table 37: Output fields: LDP session.....	380
Table 38: Output fields: LDP status.....	384
Table 39: Output fields: LDP targeted peers.....	387
Table 40: Base PCEP TLVs, objects, and messages.....	400
Table 41: PCEP path profile extension objects and TLVs.....	413
Table 42: Output fields: PCEP PCC.....	421
Table 43: Output fields: LSP.....	423
Table 44: Output fields: Path request.....	425

Table 45: Output fields: PCC peer.....426

Table 46: Output fields: PCC status.....428

List of figures

Figure 1: Label placement.....	24
Figure 2: Label packet placement.....	25
Figure 3: Bypass tunnel nodes.....	28
Figure 4: FRR node-protection example.....	30
Figure 5: MPLS-TP model.....	33
Figure 6: MPLS-TP provider edge and gateway, VLL services.....	34
Figure 7: MPLS-TP provider edge and gateway, spoke-SDP termination on VPLS.....	34
Figure 8: MPLS-TP LSR.....	35
Figure 9: MPLS-TP maintenance architecture.....	36
Figure 10: MPLS-TP LSP and tunnel information model.....	37
Figure 11: MPLS-TP PW information model.....	38
Figure 12: Example usage of PW identifiers.....	39
Figure 13: Label for LSP and PW G-ACh packets.....	40
Figure 14: BFD used for proactive CC on MPLS-TP LSP.....	42
Figure 15: BFD used for proactive CV on MPLS-TP LSP.....	42
Figure 16: Normal operation.....	45
Figure 17: Failed condition.....	46
Figure 18: Failed condition - switching at A.....	46
Figure 19: Failed condition - switching at Z.....	47
Figure 20: Establishing LSPs.....	57
Figure 21: LSP using RSVP path set up.....	58

Figure 22: Shared Risk Link Groups.....	73
Figure 23: Automatic ABR node selection for inter-area LSP.....	75
Figure 24: CSPF for an inter-area LSP.....	76
Figure 25: ABR protection using dynamic bypass LSP.....	78
Figure 26: Application of P2MP LSP in video broadcast.....	80
Figure 27: MPLS and RSVP configuration and implementation flow.....	86
Figure 28: Manual bypass tunnels.....	93
Figure 29: Subsystem interrelationships.....	257
Figure 30: LDP adjacency and session over unnumbered interface.....	264
Figure 31: LDP over RSVP application.....	266
Figure 32: LDP over RSVP application variant.....	267
Figure 33: Topology with primary and LFA routes.....	275
Figure 34: Example topology with broadcast interfaces.....	276
Figure 35: LDP adjacency and session over an IPv6 interface.....	278
Figure 36: LDP IPv6 address and FEC distribution procedure.....	284
Figure 37: LDP IPv6 address and FEC distribution procedure.....	285
Figure 38: FEC resolution in LAN.....	286
Figure 39: Basic LDP parameter provisioning.....	292
Figure 40: LDP configuration and implementation.....	292
Figure 41: NSP functional modules.....	398
Figure 42: NRC-P architecture.....	399
Figure 43: PCEP session initialization.....	402

1 Getting started

This chapter provides process flow information to configure MPLS, RSVP, and LDP protocols, as well as an overview of the document organization, 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 services and protocol support provided by the following 7210 SAS platforms, operating in one of the modes described in [Table 1: Supported modes of operation and configuration methods](#). If multiple modes of operation apply, they are explicitly noted in the topic:

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

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



Note:

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

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

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

1.1.1 Document structure and content

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

The following table lists the tasks necessary to configure MPLS applications functions.

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

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

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

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

Table 4: Configuration process

Area	Task	Chapter/section
Protocol configuration	Configure MPLS protocols:	
	• MPLS	MPLS
	• RSVP	RSVP
	• LDP	Label Distribution Protocol
	• PCEP	PCEP
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 MPLS and RSVP

This chapter provides information to configure MPLS and RSVP.

2.1 MPLS

Multiprotocol Label Switching (MPLS) is a label switching technology that provides the ability to set up connection-oriented paths over a connectionless IP network. MPLS facilitates network traffic flow and provides a mechanism to engineer network traffic patterns independently from routing tables. MPLS sets up a specific path for a sequence of packets. The packets are identified by a label inserted into each packet. MPLS is not enabled by default and must be explicitly enabled.

MPLS is independent of any routing protocol but is considered multiprotocol because it works with the Internet Protocol (IP) and frame relay network protocols.

The 7210 SAS routers enable service providers to deliver virtual private networks (VPNs) and Internet access using MPLS tunnels, with Ethernet interfaces.

2.1.1 MPLS label stack

MPLS requires a set of procedures to enhance network layer packets with label stacks, which turns them into labeled packets. Routers that support MPLS are called Label Switching Routers (LSRs). To transmit a labeled packet on a specific data link, an LSR must support the encoding technique which, when given a label stack and a network layer packet, produces a labeled packet.

In MPLS, packets can carry not just one label, but a set of labels in a stack. An LSR can swap the label at the top of the stack, pop the stack, or swap the label and push one or more labels into the stack. The processing of a labeled packet is completely independent of the level of hierarchy. The processing is always based on the top label, without regard for the possibility that some number of other labels may have been above it in the past, or that some number of other labels may be below it at present.

As described in RFC 3032, *MPLS Label Stack Encoding*, the label stack is represented as a sequence of label stack entries. Each label stack entry is represented by 4 octets. The following figure shows the label placement in a packet. [Table 5: Packet/label field description](#) describes packet and label fields.

Figure 1: Label placement





Table 5: Packet/label field description

Field	Description
Label	This 20-bit field carries the actual value (unstructured) of the label.
Exp	This 3-bit field is reserved for experimental use. It is currently used for Class of Service (CoS).
S	This bit is set to 1 for the last entry (bottom) in the label stack, and 0 for all other label stack entries.
TTL	This 8-bit field is used to encode a TTL value.

A stack can carry several labels, organized in a last in/first out order. The top of the label stack appears first in the packet and the bottom of the stack appears last, as shown in the following figure.

Figure 2: Label packet placement



OSSG014

The label value at the top of the stack is looked up when a labeled packet is received. A successful lookup reveals:

- The next hop where the packet is to be forwarded.
- The operation to be performed on the label stack before forwarding.

In addition, the lookup may reveal outgoing data link encapsulation and other information needed to properly forward the packet.

An empty label stack can be thought of as an unlabeled packet. An empty label stack has zero (0) depth. The label at the bottom of the stack is referred to as the Level 1 label. The label above it (if it exists) is the Level 2 label, and so on. The label at the top of the stack is referred to as the Level m label.

Labeled packet processing is independent of the level of hierarchy. Processing is always based on the top label in the stack which includes information about the operations to perform on the packet's label stack.

2.1.1.1 Label values

Packets traveling along an LSP (see [Label Switching Routers](#)) are identified by its label, the 20-bit, unsigned integer. The range is 0 through 1,048,575. Label values 0-15 are reserved and are defined below as follows:

- A value of 0 represents the IPv4 Explicit NULL Label. This Label value is legal only at the bottom of the Label stack. It indicates that the Label stack must be popped, and the packet forwarding must be based on the IPv4 header.
- A value of 1 represents the router alert Label. This Label value is legal anywhere in the Label stack except at the bottom. When a received packet contains this Label value at the top of the Label stack, it is delivered to a local software module for processing. The actual packet forwarding is determined by the Label beneath it in the stack. However, if the packet is further forwarded, the router alert Label should be pushed back onto the Label stack before forwarding. The use of this Label is analogous to

the use of the router alert option in IP packets. Because this Label cannot occur at the bottom of the stack, it is not associated with a particular network layer protocol.

- A value of 3 represents the Implicit NULL Label. This is a Label that a Label Switching Router (LSR) can assign and distribute, but which never actually appears in the encapsulation. When an LSR would otherwise replace the Label at the top of the stack with a new Label, but the new Label is Implicit NULL, the LSR pops the stack instead of doing the replacement. Although this value may never appear in the encapsulation, it needs to be specified in the RSVP, so a value is reserved.
- Values 4-15 are reserved for future use.

7210 SAS devices uses labels for MPLS, RSVP-TE, and LDP, as well as packet-based services such as VLL and VPLS.

Label values 16 through 1,048,575 are defined as follows:

- Label values 16 through 31 are reserved for future use.
- Label values 32 through 1,023 are available for static assignment.
- Label values 1,024 through 2,047 are reserved for future use.
- Label values 2,048 through 18,431 are statically assigned for services.
- Label values 32768 through 131,071 are dynamically assigned for both MPLS and services.
- Label values 131,072 through 1,048,575 are reserved for future use.

2.1.2 Label Switching Routers

LSRs perform the label switching function. LSRs perform different functions based on its position in an LSP. Routers in an LSP do one of the following:

- The router at the beginning of an LSP is the ingress label edge router (ILER). The ingress router can encapsulate packets with an MPLS header and forward it to the next router along the path. An LSP can only have one ingress router.
- A Label Switching Router (LSR) can be any intermediate router in the LSP between the ingress and egress routers. An LSR swaps the incoming label with the outgoing MPLS label and forwards the MPLS packets it receives to the next router in the MPLS path (LSP). An LSP can have 0-253 transit routers.
- The router at the end of an LSP is the egress label edge router (ELER). The egress router strips the MPLS encapsulation which changes it from an MPLS packet to a data packet, and then forwards the packet to its final destination using information in the forwarding table. Each LSP can have only one egress router. The ingress and egress routers in an LSP cannot be the same router.

A router in your network can act as an ingress, egress, or transit router for one or more LSPs, depending on your network design.

An LSP is confined to one IGP area for LSPs using constrained-path. They cannot cross an autonomous system (AS) boundary.

Static LSPs can cross AS boundaries. The intermediate hops are manually configured so the LSP has no dependence on the IGP topology or a local forwarding table.

2.1.2.1 LSP types

The following are LSP types:

- **static LSPs**

A static LSP specifies a static path. All routers that the LSP traverses must be configured manually with labels. No signaling such as RSVP or LDP is required.

- **signaled LSP**

LSPs are set up using a signaling protocol such as RSVP-TE or LDP. The 7210 SAS supports only RSVP-TE for setting up LSPs. The signaling protocol allows labels to be assigned from an ingress router to the egress router. Signaling is triggered by the ingress routers. Configuration is required only on the ingress router and is not required on intermediate routers. Signaling also facilitates path selection.

There are two signaled LSP types:

- **explicit-path LSPs**

MPLS uses RSVP-TE to set up explicit path LSPs. The hops within the LSP are configured manually. The intermediate hops must be configured as either strict or loose meaning that the LSP must take either a direct path from the previous hop router to this router (strict) or can traverse through other routers (loose). You can control how the path is set up. They are similar to static LSPs but require less configuration. See [RSVP](#).

- **constrained-path LSPs**

The intermediate hops of the LSP are dynamically assigned. A constrained path LSP relies on the Constrained Shortest Path First (CSPF) routing algorithm to find a path which satisfies the constraints for the LSP. In turn, CSPF relies on the topology database provided by the extended IGP such as OSPF or IS-IS.

When the path is found by CSPF, RSVP uses the path to request the LSP set up. CSPF calculates the shortest path based on the constraints provided such as bandwidth, class of service, and specified hops.

If fast reroute is configured, the ingress router signals the routers downstream. Each downstream router sets up a detour for the LSP. If a downstream router does not support fast reroute, the request is ignored and the router continues to support the LSP. This can cause some of the detours to fail, but otherwise the LSP is not impacted.

No bandwidth is reserved for the rerouted path. If the user enters a value in the bandwidth parameter in the **config>router>mpls>lsp>fast-reroute** context, it will have no effect on the LSP backup LSP establishment.

Hop-limit parameters specifies the maximum number of hops that an LSP can traverse, including the ingress and egress routers. An LSP is not set up if the hop limit is exceeded. The hop count is set to 255 by default for the primary and secondary paths. It is set to 16 by default for a bypass or detour LSP path.

2.1.2.2 MPLS Fast Re-Route (FRR)

The MPLS facility bypass method of MPLS Fast Re-Route (FRR) functionality is extended to the ingress node.

The behavior of an LSP at an ingress LER with both fast reroute and a standby LSP path configured is as follows:

- When a downstream detour becomes active at a point of local repair (PLR):

The ingress LER switches to the standby LSP path. If the primary LSP path is repaired subsequently at the PLR, the LSP will switch back to the primary path. If the standby goes down, the LSP is switched

back to the primary, even though it is still on the detour at the PLR. If the primary goes down at the ingress while the LSP is on the standby, the detour at the ingress is cleaned up and for one-to-one detours a "path tear" is sent for the detour path. In other words, the detour at the ingress does not protect the standby. If and when the primary LSP is again successfully re-signaled, the ingress detour state machine will be restarted.

- When the primary fails at the ingress:

The LSP switches to the detour path. If a standby is available then LSP would switch to standby on expiration of **hold-timer**. If **hold-timer** is disabled then switchover to standby would happen immediately. On successful global revert of primary path, the LSP would switch back to the primary path.

- Admin groups are not taken into account when creating detours for LSPs.

2.1.2.3 Manual bypass LSP

The 7210 SAS supports Manual bypass tunnels, on implementation of the Manual bypass feature a LSP can be preconfigured from a PLR which is used exclusively for bypass protection. If a path message for a new LSP requests for bypass protection, the node checks if a manual bypass tunnel satisfying the path constraints exists. If a tunnel is found, it is selected. If no such tunnel exists by default, the 7210 SAS dynamically signals a bypass LSP.


Users can disable the dynamic bypass creation on a per node basis using the CLI.

A maximum of 1000 associations of primary LSP paths can be made with a single manual bypass at the PLR node. If dynamic bypass creation is disabled on the node, it is recommended to configure additional manual bypass LSPs to handle the required number of associations.

2.1.2.3.1 PLR bypass LSP selection rules

The following figure shows bypass tunnel nodes.

Figure 3: Bypass tunnel nodes



The PLR uses the following rules to select a bypass LSP among multiple manual and dynamic bypass LSPs at the time of establishment of the primary LSP path or when searching for a bypass for a protected LSP which does not have an association with a bypass tunnel:

1. The MPLS/RSVP task in the PLR node checks if an existing manual bypass satisfies the constraints. If the path message for the primary LSP path indicated node protection is needed, which is the default LSP FRR setting at the head end node, MPLS/RSVP task searches for a node-protect' bypass LSP. If the path message for the primary LSP path indicated link protection is needed, then it searches for a link-protect bypass LSP.

2. If multiple manual bypass LSPs satisfying the path constraints exist, it will prefer a manual-bypass terminating closer to the PLR over a manual bypass terminating further away. If multiple manual bypass LSPs satisfying the path constraints terminate on the same downstream node, it selects one with the lowest IGP path cost or if in a tie, picks the first one available.
3. If none satisfies the constraints and dynamic bypass tunnels have not been disabled on PLR node, then the MPLS/RSVP task in the PLR will check if any of the already established dynamic bypasses of the requested type satisfies the constraints.
4. If none do, then the MPLS/RSVP task will ask CSPF to check if a new dynamic bypass of the requested type, node-protect or link-protect, can be established.
5. If the path message for the primary LSP path indicated node protection is needed, and no manual bypass was found after Step 1, or no dynamic bypass LSP was found after 3 attempts of performing Step 3, the MPLS/RSVP task will repeat Steps 1-3 looking for a suitable link-protect bypass LSP. If none are found, the primary LSP will have no protection and the PLR node must clear the "local protection available" flag in the IPv4 address sub-object of the record route object (RRO) starting in the next Resv refresh message it sends upstream.
6. If the path message for the primary LSP path indicated link protection is needed, and no manual bypass was found after step 1, or no dynamic bypass LSP was found after performing Step 3, the primary LSP will have no protection and the PLR node must clear the "local protection available" flag in the IPv4 address sub-object of the RRO starting in the next RESV refresh message it sends upstream. The PLR will not search for a node-protect' bypass LSP in this case.
7. If the PLR node successfully makes an association, it must set the "local protection available" flag in the IPv4 address sub-object of the RRO starting in the next RESV refresh message it sends upstream.
8. For all primary LSP that requested FRR protection but are not currently associated with a bypass tunnel, the PLR node on reception of RESV refresh on the primary LSP path repeats Steps 1-7.

If the user disables dynamic-bypass tunnels on a node while dynamic bypass tunnels were activated and were passing traffic, traffic loss will occur on the protected LSP. Furthermore, if no manual bypass exist that satisfy the constraints of the protected LSP, the LSP will remain without protection.

If the user configures a bypass tunnel on node B and dynamic bypass tunnels have been disabled, LSPs which have been previously signaled and which were not associated with any manual bypass tunnel, for example, none existed, will be associated with the manual bypass tunnel if suitable. The node checks for the availability of a suitable bypass tunnel for each of the outstanding LSPs every time a RESV message is received for these LSPs.

If the user configures a bypass tunnel on node B and dynamic bypass tunnels have not been disabled, LSPs which have been previously signaled over dynamic bypass tunnels will not automatically be switched into the manual bypass tunnel even if the manual bypass is a more optimized path. The user will have to perform a make before break at the head end of these LSPs.

If the manual bypass goes into the down state in node B and dynamic bypass tunnels have been disabled, node B (PLR) will clear the "protection available" flag in the RRO IPv4 sub-object in the next RESV refresh message for each affected LSP. It will then try to associate each of these LSPs with one of the manual bypass tunnels that are still up. If it finds one, it will make the association and set again the "protection available" flag in the next RESV refresh message for each of these LSPs. If it could not find one, it will keep checking for one every time a RESV message is received for each of the remaining LSPs. When the manual bypass tunnel is back UP, the LSPs which did not find a match will be associated back to this tunnel and the protection available flag is set starting in the next RESV refresh message.

If the manual bypass goes into the down state in node B and dynamic bypass tunnels have not been disabled, node B will automatically signal a dynamic bypass to protect the LSPs if a suitable one does not exist. Similarly, if an LSP is signaled while the manual bypass is in the down state, the node will only signal


a dynamic bypass tunnel if the user has not disabled dynamic tunnels. When the manual bypass tunnel is back into the UP state, the node will not switch the protected LSPs from the dynamic bypass tunnel into the manual bypass tunnel.

2.1.2.3.2 FRR node-protection (facility)

The MPLS Fast Re-Route (FRR) functionality enables PLRs to be aware of the missing node protection and allows them to regularly probe for a node-bypass.

The following figure shows an LSP scenario.

Figure 4: FRR node-protection example



Where:

- LSP 1: between PE_1 to PE_2, with CSPF, FRR facility node-protect enabled.
- P_1 protects P_2 with bypass-nodes P_1 - P_3 - P_4 - PE_4 - PE_2.
- If P_4 fails, P_1 tries to establish the bypass-node three times.
- When the bypass-node creation fails, P_1 will protect link P_1-P_2.
- P_1 protects the link to P_2 through P_1 - P_5 - P_2.
- P_4 returns online.

As LSP 1 had requested node protection, but because of a lack of any available path, it could only obtain link protection. Therefore, every 60 seconds the PLR for LSP 1 will search for a new path that may be able to provide node protection. When P_4 is back online and such a path is available, A new bypass tunnel will be signaled and LSP 1 will get associated with this new bypass tunnel.

2.1.2.3.3 Uniform FRR failover time

The failover time during FRR consists of a detection time and a switchover time. The detection time corresponds to the time it takes for the RSVP control plane protocol to detect that a network IP interface is down or that a neighbor/next-hop over a network IP interface is down. The control plane can be informed of an interface down event when the event is because of a failure in a lower layer such in the physical layer. The control plane can also detect the failure of a neighbor/next-hop on its own by running a protocol such as Hello, Keep-Alive, or BFD.

The switchover time is measured from the time the control plane detected the failure of the interface or neighbor/next-hop to the time the IOM completed the reprogramming of all the impacted ILM or service records in the datapath. This includes the time it takes for the control plane to send a down notification to all IOMs to request a switch to the backup NHLFE.

Uniform Fast-Reroute (FRR) failover enables the switchover of MPLS and service packets from the outgoing interface of the primary LSP path to that of the FRR backup LSP within the same amount of time regardless of the number of LSPs or service records. This is achieved by updating Ingress Label Map (ILM) records and service records to point to the backup Next-Hop Label to Forwarding Entry (NHLFE) in a single operation.

2.1.3 Configuration guidelines

Implicit NULL must be enabled for use of Manual Bypass or Dynamic Bypass (FRR facility) if the 7210 is used as a egress LER or is a Merge Point.

2.2 MPLS pseudowire hash label support

The MPLS pseudowire hash label allows LSR nodes in a network to load balance labeled packets in a much more granular fashion than allowed by simply hashing on the standard label stack. It also removes the need to have an LSR inspect the payload below the label stack to check for an IPv4 or IPv6 header.

An MPLS hash label, is inserted by the ingress LER at the bottom of the label stack in packets forwarded over an LSP. The value of the label is the result of the hash of the packet headers (the packet header fields used depends on the capability of the ingress LER node). Since the ingress LER hash routine maintains packet ordering within a conversation, this guarantees that the spraying of packets by an LSR hashing on the extended label stack, which includes the hash label, will also maintain packet ordering within a conversation. LSR hashing pertains to multiple LDP ECMP paths or multiple paths over a LAG network port.



Note:

- On 7210 SAS devices, the ingress node does not use the pseudowire hash label for ECMP hashing and LAG hashing. It is only available for use by the transit MPLS LSR nodes. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide* for information about the fields used by the ingress LER for ECMP and LAG hashing.
- The pseudowire hash label is supported for VLL with spoke SDP, and VPLS services with spoke SDP and mesh SDP.
- This feature is only supported on the 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-R6 (IMMv2 cards), and 7210 SAS-R12 (IMMv2 cards).
- The pseudowire hash label is not accounted for in the total number of MPLS transport and service labels that the node pushes and pops.

2.3 MPLS Transport Profile (MPLS-TP)



Note:

This feature is only supported on 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T.

MPLS can be used to provide a network layer to support packet transport services. In some operational environments, it is desirable that the operation and maintenance of such an MPLS based packet transport network follow operational models typical in traditional optical transport networks (For example, SONET/SDH), while providing additional OAM, survivability and other maintenance functions targeted at that environment.

MPLS-TP defines a profile of MPLS targeted at transport applications. This profile defines the specific MPLS characteristics and extensions required to meet transport requirements, while retaining compliance to the standard IETF MPLS architecture and label switching paradigm. The basic requirements architecture for MPLS-TP are described by the IETF in RFC 5654, RFC 5921 and RFC 5960, to meet two objectives:

- To enable MPLS to be deployed in a transport network and operated in a similar manner to existing transport technologies.
- To enable MPLS to support packet transport services with a similar degree of predictability to that found in existing transport networks.

To meet these objectives, MPLS-TP has a number of high level characteristics:

- It does not modify the MPLS forwarding architecture, which is based on existing pseudowire and LSP constructs. Point-to-point LSPs may be unidirectional or bidirectional. Bi-directional LSPs must be congruent (that is, co-routed and follow the same path in each direction). The 7210 SAS supports bidirectional co-routed MPLS-TP LSPs.
- There is no LSP merging.
- OAM, protection and forwarding of data packets can operate without IP forwarding support. When static provisioning is used, there is no dependency on dynamic routing or signaling.
- LSP and pseudowire monitoring is only achieved through the use of OAM and does not rely on control plane or routing functions to determine the health of a path. For example, LDP hello failures, do not trigger protection.
- MPLS-TP can operate in the absence of an IP control plane and IP forwarding of OAM traffic. In 7210 SAS releases, MPLS-TP is only supported on static LSPs and PWs.

The 7210 SAS supports MPLS-TP on LSPs and PWs with static labels. MPLS-TP is not supported on dynamically signaled LSPs and PWs. MPLS-TP is supported for Epipe, and Epipe Spoke SDP termination on VPLS. Static PWs may use SDPs that use either static MPLS-TP LSPs or RSVP-TE LSPs.

The following MPLS-TP OAM and protection mechanisms, defined by the IETF, are supported:

- MPLS-TP Generic Associated Channel for LSPs and PWs (RFC 5586)
- MPLS-TP Identifiers (RFC 6370)
- Proactive CC, CV, and RDI using BFD for LSPs (RFC 6428)
- BFD based CV is not supported in this release.
- On-Demand CV for LSPs and PWs using LSP Ping and LSP Trace (RFC 6426)
- 1-for-1 Linear protection for LSPs (RFC 6378)


- Static PW Status Signaling (RFC 6478)

The 7210 SAS can play the role of an LER and an LSR for static MPLS-TP LSPs, and a PE/T-PE for static MPLS-TP PWs. It can also act an MPLS network that supports both MPLS-TP and dynamic IP/MPLS.

2.3.1 MPLS-TP model

The following figure shows a high level functional model for MPLS-TP in 7210 SAS. LSP A and LSP B are the working and protect LSPs of an LSP tunnel. These are modeled as working and protect paths of an MPLS-TP LSP in 7210 SAS. MPLS-TP OAM runs in-band on each path. 1:1 linear protection coordinates the working and protect paths, using a protection switching coordination protocol (PSC) that runs in-band on each path over a Generic Associated Channel (G-ACh) on each path. Each path can use an IP numbered, IP unnumbered, or MPLS-TP unnumbered (that is, non-IP) interface.

Figure 5: MPLS-TP model



For 7210 SAS platforms, all MPLS-TP LSPs are bidirectional co-routed as detailed in RFC 5654. That is, the forward and backward directions follow the same route (in terms of links and nodes) across the network. Both directions are setup, monitored and protected as a single entity. Therefore, both ingress and egress directions of the same LSP segment are associated at the LER and LSR and use the same interface (although this is not enforced by the system).

In the above model, an SDP can use one MPLS-TP LSP. This abstracts the underlying paths toward the overlying services, which are transported on pseudowires. Pseudowires are modeled as spoke SDPs and can also use MPLS-TP OAM. PWs with static labels may use SDPs that in-turn use either signaled RSVP-TE LSPs, or one static MPLS-TP LSP.

2.3.2 MPLS-TP provider edge and gateway

This section describes examples of roles for the 7210 SAS in an MPLS-TP network.

2.3.2.1 VLL services

The 7210 SAS may use MPLS TP LSPs, and PWs, to transport point to point virtual leased line services. The node plays the role of a terminating PE or switching PE for VLLs. Only T-PE functionality is supported on 7210 SAS-T (network mode). Both T-PE and S-PE (static only) functionality is supported on 7210 SAS-R6 and 7210 SAS-R12. Epipe is supported.

The following figures show the use of the 7210 SAS as a T-PE/S-PE for services in an MPLS-TP domain.

Figure 6: MPLS-TP provider edge and gateway, VLL services





Figure 7: MPLS-TP provider edge and gateway, spoke-SDP termination on VPLS




Note:

- Spoke SDP termination on IES and VPRN services is supported on all 7210 SAS platforms as described in this document, except the 7210 SAS-R6 and 7210 SAS-R12. Is it also supported on 7210 SAS platforms operating in access-uplink mode.
- MPLS-TP Epipe spoke SDP termination on VPLS is supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

- 7210 SAS-R6 and 7210 SAS-R12 nodes can be used as T-PE nodes (LER nodes) that originate Epipe services, or as S-PE nodes acting as stitching points for MPLS-TP PWs.

Figure 8: MPLS-TP LSR



2.3.3 Detailed descriptions of MPLS-TP

2.3.3.1 MPLS-TP LSPs

The 7210 SAS supports the configuration of MPLS-TP tunnels, which comprise a working and, optionally, a protect LSP. In 7210 SAS, a tunnel is referred to as an LSP, while an MPLS-TP LSP is referred to as a path. It is then possible to bind an MPLS-TP tunnel to an SDP.

MPLS-TP LSPs (that is, paths) with static labels are supported. MPLS-TP is not supported for signaled LSPs.

Both bidirectional associated (where the forward and reverse directions of a bidirectional LSP are associated at a specific LER, but may take different routes through the intervening network) and bidirectional co-routed (where the forward and reverse directions of the LSP are associated at each LSR, and take the same route through the network) are possible in MPLS-TP. However, only bidirectional co-routed LSPs are supported.

It is possible to configure MPLS-TP identifiers associated with the LSP, and MPLS-TP OAM parameters on each LSP of a tunnel. MPLS-TP protection is configured for a tunnel at the level of the protect path level. Both protection and OAM configuration is managed through templates to simplify provisioning for a large numbers of tunnels.

The 7210 SAS plays the role of either an LER or an LSR.

2.3.3.2 MPLS-TP on pseudowires

MPLS-TP is supported on PWs with static labels. The provisioning model supports RFC 6370-style PW path identifiers for MPLS-TP PWs.

MPLS-TP PWs reuse the static PW provisioning model of previous 7210 SAS releases. The primary distinguishing feature for an "MPLS-TP" PW is the ability to configure MPLS-TP PW path identifiers, and to support MPLS-TP OAM and static PW status signaling.

The 7210 SAS can perform the role of a T-PE for a PW with MPLS-TP. The 7210 SAS-R6 and 7210 SAS-R12 can perform the role of S-PE for MPLS-TP PW. The 7210 SAS-T does not support S-PE functionality.

A spoke-SDP with static PW labels and MPLS-TP identifiers and OAM capabilities can use an SDP that uses either an MPLS-TP tunnel, or that uses regular RSVP-TE LSPs. The control word is supported for all MPLS-TP PWs.


2.3.4 MPLS-TP maintenance identifiers

MPLS-TP is designed for use both with, and without, a control plane. MPLS-TP therefore specifies a set of identifiers that can be used for objects in either environment. This includes a path and maintenance identifier architecture comprising Node, Interface, PW and LSP identifiers, Maintenance Entity Groups (MEGs), Maintenance End Points (MEPs) and Maintenance Intermediate Points (MIPs). These identifiers are specified in RFC 6370.

MPLS-TP OAM and protection switching operates within a framework that is designed to be similar to existing transport network maintenance architectures. MPLS-TP introduces concept of maintenance domains to be managed and monitored. In these, Maintenance Entity Group End Points (MEPs) are edges of a maintenance domain. OAM of a maintenance level must not leak beyond corresponding MEP and so MEPs typically reside at the end points of LSPs and PWs. Maintenance Intermediate Points (MIPs) define intermediate nodes to be monitored. Maintenance Entity Groups (MEGs) comprise all the MEPs and MIPs on an LSP or PW.

The following figure shows the MPLS-TP maintenance architecture.

Figure 9: MPLS-TP maintenance architecture



al_0226

Both IP-compatible and ICC (ITU-T carrier code) based identifiers for the above objects are specified in the IETF, but only the IP-compatible identifiers defined in RFC 6370 are supported.

The 7210 SAS supports the configuration of the following node and interface related identifiers:

- **Global_ID:** In MPLS-TP, the Global_ID should be set to the AS# of the node. If not explicitly configured, then it assumes the default value of 0. In 7210 SAS, the source Global ID for an MPLS-TP Tunnel is


- taken to be the Global ID configured at the LER. The destination Global ID is optional in the tunnel configuration. If it is not configured, then it is taken as the same as the source Global ID.
- Node_ID: This is a 32-bit value assigned by the operator within the scope of the Global_ID. The 7210 SAS supports the configuration of an IPv4 formatted address <a.b.c.d> or an unsigned 32-bit integer for the MPLS-TP Node ID at each node. The node ID must be unique within the scope of the global ID, but there is no requirement for it to be a valid IP address. Indeed, a node-id can represent a separate IP-compatible addressing space that may be separate from the IP addressing plan of the underlying network. If no node ID is configured, then the node ID is taken to be the system interface IPv4 address of the node. When configuring a tunnel at an LER, either an IPv4 or an unsigned integer Node ID can be configured as the source and destination identifiers, but both ends must be of the same type.

Statically configured LSPs are identified using GMPLS-compatible identifiers with the addition of a Tunnel_Num and LSP_Num. As in RSVP-TE, tunnels represent, for example, a set of working and protect LSPs. These are GMPLS-compatible because GMPLS chosen by the IETF as the control plane for MPLS-TP LSPs, although this is not supported in 7210 SAS 6.1R1 release. PWs are identified using a PW Path ID which has the same structure as FEC129 All Type 2.

The 7210 SAS derives the identifiers for MEPs and MIPs on LSPs and PWs based on the configured identifiers for the MPLS-TP Tunnel, LSP or PW Path ID, for use in MPLS-TP OAM and protection switching, as per RFC 6370.

The information models for LSPs and PWs supported in 7210 SAS are shown in [Figure 10: MPLS-TP LSP and tunnel information model](#) and [Figure 11: MPLS-TP PW information model](#). The figures use the terminology defined in RFC 6370.

Figure 10: MPLS-TP LSP and tunnel information model




The MPLS-TP Tunnel ID and LSP ID are not to be confused with the RSVP-TE tunnel ID implemented on the 7210 system. The following table describes how these map to the X and Y ends of the tunnel shown in the preceding figure for the case of co-routed bidirectional LSPs.

Table 6: Mapping from RSVP-TE to MPLS-TP maintenance identifiers

RSVP-TE Identifier	MPLS-TP Maintenance Identifier
Tunnel Endpoint Address	Node ID (Y)
Tunnel ID (X)	Tunnel Num (X)
Extended Tunnel ID	Node ID (X)
Tunnel Sender Address	Node ID (X)
LSP ID	LSP Num

Figure 11: MPLS-TP PW information model



al_0228


In the PW information model shown in the preceding figure, the MS-PW is identified by the PW Path ID that is composed of the full AGI:SAII:TAII. The PW Path ID is also the MEP ID at the T-Pes, so a user does not have to explicitly configure a MEP ID, it is automatically derived by the system. For MPLS-TP PWs with static labels, although the PW is not signaled end-to-end, the directionality of the SAII and TAIL is taken to be the same as for the equivalent label mapping message that is, from downstream to upstream. This is to maintain consistency with signaled pseudowires using FEC 129.

On the 7210 SAS, an S-PE for an MS-PW with static labels is configured as a pair of spoke SDPs bound together in an VLL service using the vc-switching command. Therefore, the PW Path ID configured at the spoke-sdp level at an S-PE must contain the Global-ID, Node-ID and AC-ID at the far end T-Pes, not the local S-PE. The ordering of the SAII:TAII in the PW Path ID where static PWs are used should be consistent with the direction of signaling of the egress label to a spoke-SDP forming that segment, if that

label were signaled using T-LDP (in downstream unsolicited mode). VCCV Ping will check the PW ID in the VCCV Ping echo request message against the configured PW Path ID for the egress PW segment.

The following figure shows an example of how the PW Path IDs can be configured for a simple two-segment MS-PW.

Figure 12: Example usage of PW identifiers




2.3.4.1 Generic Associated Channel

MPLS-TP requires that all OAM traffic be carried in-band on both directions of an LSP or PW. This is to ensure that OAM traffic always shares fate with user data traffic. This is achieved by using an associated control channel on an LSP or PW, similar to that used today on PWs. This creates a channel, which is used for OAM, protection switching protocols (for example, LSP linear protection switching coordination), and other maintenance traffic, and is known as the Generic Associated Channel (G-ACh).

RFC 5586 specifies mechanisms for implementing the G-ACh, relying on the combination of a reserved MPLS label, the 'Generic-ACH Label (GAL)', as an alert mechanism (value=13) and Generic Associated Channel Header (G-ACh) for MPLS LSPs, and using the Generic Associated Channel Header, only, for MPLS PWs (although the GAL is allowed on PWs).

The purpose of the GAL is to indicate that a G-ACh resides at the bottom of the label stack, and is only visible when the bottom non-reserved label is popped. The G-ACh channel type is used to indicate the packet type carried on the G-ACh. Packets on a G-ACh are targeted to a node containing a MEP by ensuring that the GAL is pushed immediately below the label that is popped at the MEP (for example, LSP endpoint or PW endpoint), so that it can be inspected as soon as the label is popped. A G-ACh packet is targeted to a node containing a MIP by setting the TTL of the LSP or PW label, as applicable, so that it expires at that node, in a similar manner to the SR OS implementation of VCCV for MS-PWs.

The following figure shows labels for LSP and PW G-ACh packets.

Figure 13: Label for LSP and PW G-ACh packets

The 7210 SAS supports the G-ACh on static pseudowires and static LSPs.

2.3.4.2 MPLS-TP Operations, Administration and Maintenance (OAM)

This section details the MPLS-TP OAM mechanisms that are supported.

2.3.4.2.1 On-demand connectivity verification (CV) using LSP-Ping

MPLS-TP supports mechanisms for on demand CC/CV as well as route tracing for LSPs and PWs. These are required to enable an operator to test the initial configuration of a transport path, or to assist with fault isolation and diagnosis. On demand CC/CV and route tracing for MPLS-TP is based on LSP-Ping and is described in RFC 6426. Three possible encapsulations are specified in that RFC:

- IP encapsulation, using the same label stack as RFC 4379, or encapsulated in the IPv4 G-ACh channel with a GAL/ACH
- A non-IP encapsulation with GAL/ACH for LSPs and ACH for PWs.

In IP-encapsulation, LSP-Ping packets are sent over the MPLS LSP for which OAM is being performed and contain an IP/UDP packet within them. The On-demand CV echo response message is sent on the reverse path of the LSP, and the reply contains IP/UDP headers followed by the On-demand CV payload.

In non-IP environments, LSP ping can be encapsulated with no IP/UDP headers in a G-ACh and use a source address TLV to identify the source node, using forward and reverse LSP or PW associated channels on the same LSP or PW for the echo request and reply packets. In this case, no IP/UDP headers are included in the LSP-Ping packets.

The 7210 SAS supports the following encapsulations:

- IP encapsulation with ACH for PWs (as per VCCV type 1).
- IP encapsulation without ACH for LSPs using labeled encapsulation
- Non-IP encapsulation with ACH for both PWs and LSPs.

LSP Ping and VCCV Ping for MPLS-TP use two new FEC sub-types in the target FEC stack to identify the static LSP or static PW being checked. These are the Static LSP FEC sub-type, which has the same format as the LSP identifier described above, and the Static PW FEC sub-type. These are used in-place of the currently defined target FEC stack sub-TLVs.

In addition, MPLS-TP uses a source/destination TLV to carry the MPLS-TP global-id and node-id of the target node for the LSP ping packet, and the source node of the LSP ping packet.

LSP Ping and VCCV-Ping for MPLS-TP can only be launched by the LER or T-PE. The replying node therefore sets the TTL of the LSP label or PW label in the reply packet to 255 to ensure that it reaches the node that launched the LSP ping or VCCV Ping request.

Downstream mapping support

The following table describes the four RFC 4379 specified address types for the downstream mapping TLV for use with IP numbered and unnumbered interfaces.

Table 7: Address types for the downstream mapping TLV

Type #	Address type	K octets	Reference	7210 SAS-R6 and 7210 SAS-R12 support	7210 SAS-T support
1	IPv4 Numbered	16	RFC 4379	Yes	Yes
2	IPv4 Unnumbered	16	RFC 4379	Yes	Yes
3	IPv6 Numbered	40	RFC 4379	No	No
4	IPv6 Unnumbered	28	RFC 4379	No	No

RFC 6426 adds address type 5 for use with Non IP interfaces, including MPLS-TP interfaces. In addition, this RFC specifies that type 5 must be used when non-IP ACH encapsulation is used for LSP Trace.

It is possible to send and respond to a DSMAP/DDMAP TLV in the LSP Trace packet for numbered IP interfaces as per RFC 4379. In this case, the echo request message contains a downstream mapping TLV with address type 1 (IPv4 address) and the IPv4 address in the DDMAP/DSMAP TLV is taken to be the IP address of the IP interface that the LSP uses. The LSP trace packet therefore contains a DSMAP TLV in addition to the MPLS-TP static LSP TLV in the target FEC stack.

DSMAP/DDMAP is not supported for pseudowires.

2.3.4.2.2 Proactive CC, CV, and RDI

Proactive Continuity Check (CC) is used to detect a loss of continuity defect (LOC) between two MEPs in a MEG. Proactive Connectivity Verification (CV) is used to detect an unexpected connectivity defect between two MEPs (For example: mis-merging or disconnection), as well as unexpected connectivity within the MEG with an unexpected MEP. This feature implements both functions using proactive generation of OAM packets by the source MEP that are processed by the peer sink MEP. CC and CV packets are always sent in-band such that they share with user traffic, either on an LSP, PW or section and are used to trigger protection switching mechanisms.

Proactive CC/CV based on bidirectional forwarding detection (BFD) for MPLS-TP is described in RFC 6428. BFD packets are sent using operator configurable timers and encapsulated without UDP/IP headers on a standardized G-ACh channel on an LSP or PW.

CC packets simply consist of a BFD control packet, while CV packets also include an identifier for the source MEP in order that the sink MEP can detect if it is receiving packets from an incorrect peer MEP, therefore indicating a mis-connectivity defect. Other defect types (including period mis-configuration defect) should be supported. When a supported defect is detected, an appropriate alarm is generated (for

example: log, SNMP trap) at the receiving MEP and all traffic on the associated transport path (LSP or PW) is blocked. This is achieved using linear protection for CC defects, and by blocking the ingress datapath for CV defects.

**Note:**

7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T only support BFD-based CC mode. BFD-based CC-CV mode is not supported in the current release.

When an LSP with CV is first configured, the LSP will be held in the CV defect state for 3.5 seconds after the first valid CV packet is received.

The following figures show an example of linear protection switching of LSPs triggered based on a CC or CV defect detected by BFD CC/CV.

Figure 14: BFD used for proactive CC on MPLS-TP LSP





Figure 15: BFD used for proactive CV on MPLS-TP LSP



Note that RFC 6428 defines two BFD session modes: Coordinated mode, in which the session state on both directions of the LSP is coordinated and constructed from a single, bidirectional BFD session, and independent mode, in which two independent sessions are bound together at a MEP. Coordinated mode is supported.

BFD is supported on MPLS-TP LSPs. When BFD_CV detects a mis-connectivity on an LSP, the system will drop all incoming non-OAM traffic with the LSP label (at the LSP termination point) instead of forwarding it to the associated SAP or PW segment.

The following G-ACh channel types are supported for the combined CC/CV mode:

- 0x22 for BFD CC with no IP encapsulation
- 0x23 for BFD CV

The 0x07 G-ACh channel type is used for the CC-only mode.

2.3.4.2.3 BFD-based RDI

RDI provides a mechanism whereby the source MEP can be informed of a downstream failure on an LSP, and can therefore either raise an alarm, or initiate a protection switching operation. In the case of BFD based CC/CV, RDI is communicated using the BFD diagnostic field in BFC CC/CV messages. The following diagnostic codes are supported:

- 1 - Control Detection Time Expired
- 9 - mis-connectivity defect

2.3.4.3 PW control channel status notifications (static pseudowire status signaling)

MPLS-TP introduces the ability to support a full range of OAM and protection / redundancy on PWs for which no dynamic T-LDP control plane exists. Static PW status signaling is used to advertise the status of a PW with statically configured labels by encapsulating the PW status TLV in a G-ACh on the PW. This mechanism enables OAM message mapping and PW redundancy for such PWs, as defined in RFC6478. This mechanism is known as control channel status signaling in SR OS.

PW control channel status notifications use a similar model to T-LDP status signaling. That is, in general, status is always sent to the nearest neighbor T-PE. To achieve this, the PW label TTL is set to 1 for the G-ACh packet containing the status message.

Control channel status notifications are disabled by default on a spoke-sdp. If they are enabled, then the default refresh interval is set to zero (although this value should be configurable in CLI). That is, when a status bit changes, three control channel status packets will be sent consecutively at one-second intervals, and then the transmitter will fall silent. If the refresh timer interval is non-zero, then status messages will continue to be sent at that interval. The system supports the configuration of a refresh timer of 0, or from 10-65535 seconds. The recommended value is 600 seconds.

In order to constrain the CPU resources consumed processing control channel status messages, the system implements a credit-based mechanism. If a user enables control channel status on a PW[n], then a certain number of credits *c_n* are consumed from a CPM-wide pool of *max_credit* credits. The number of credits consumed is inversely proportional to the configured refresh timer (the first three messages at 1 second interval do not count against the credit). If the *current_credit* ≤ 0, then control channel status signaling cannot be configured on a PW (but the PW can still be configured and no shutdown).

If a PE with a non-zero refresh timer configured does not receive control channel status refresh messages for 3.5 time the specified timer value, then by default it will time out and assume a PW status of zero.

A trap is generated if the refresh timer times out.

If PW redundancy is configured, the system will always consider the literal value of the PW status; a time-out of the refresh timer will not impact the choice of the active transit object for the VLL service. The result of this is that if the refresh timer times-out, and a given PW is currently the active PW, then the system will

not fail-over to an alternative PW if the status is zero and some lower-layer OAM mechanism for example, BFD has not brought down the LSP due to a connectivity defect. It is recommended that the PW refresh timer be configured with a much longer interval than any proactive OAM on the LSP tunnel, so that the tunnel can be brought down before the refresh timer expires if there is a CC defect.

A unidirectional continuity fault on a RSVP TE LSP may not result in the LSP being brought down before the received PW status refresh timer expires. It is therefore recommended that either bidirectional static MPLS-TP LSPs with BFD CC, or additional protection mechanisms. For example, FRR be used on RSVP-TE LSPs carrying MPLS-TP PWs. This is particularly important in active/standby PW dual homing configurations, where the active / standby forwarding state or operational state of every PW in the redundancy set must be accurately reflected at the redundant PE side for the configuration.

**Note:**

A PW with a refresh timer value of zero is always treated as having not expired.

The 7210 SAS implements a hold-down timer for control-channel-status pw-status bits in order to suppress bouncing of the status of a PW. For a specific spoke-sdp, if the system receives 10 pw-status "change" events in 10 seconds, the system will "hold-down" the spoke-sdp on the local node with the last received non-zero pw-status bits for 20 seconds. It will update the local spoke with the most recently received pw-status. This hold down timer is not persistent across shutdown/no-shutdown events.

2.3.4.4 Pseudowire redundancy and active / standby dual-homing

PW redundancy is supported for static MPLS-TP pseudowires. However, instead of using T-LDP status signaling to signal the forwarding state of a PW, control channel status signaling is used.

The following PW redundancy scenarios are available:

- MC-LAG with single and multi-segment PWs interconnecting the PEs.
- The 7210 SAS-T can only act as T-PE when a multi-segment PW is used.
- The 7210 SAS-R6 and 7210 SAS-R12 can act as T-PE or S-PE when a multi-segment PW is used.
- Dual-homing of a VLL service into redundant IES or VPRN PEs (the IES and VPRN service is configured on the 7750 PEs), with active/standby PWs.
 - In this scenario, 7210 SAS originates the Epipe MPLS-TP PWs as a T-PE. 7210 SAS nodes cannot terminate a MPLS-TP PW in a IES or VPRN service.
- Dual-homing of a VLL service into a VPLS with active/standby PWs.
 - In this scenario, 7210 SAS originates the Epipe MPLS-TP PWs as a T-PE. 7210 SAS-R6 and 7210 SAS-R12 nodes can terminate a MPLS-TP PW in a VPLS service.

**Note:**

Active/standby dual-homing into routed VPLS is not supported for MPLS-TP PWs.

It is possible to configure inter-chassis backup (ICB) PWs as static MPLS-TP PWs with MPLS-TP identifiers. Only MPLS-TP PWs are supported in the same endpoint. That is, PWs in an endpoint must either be all MPLS-TP, or none of them must be MPLS-TP. This implies that an ICB used in an endpoint for which other PWs are MPLS TP must also be configured as an MPLS-TP PW.

A fail over to a standby pseudowire is initiated based on the existing supported methods (For example, failure of the SDP).

2.3.4.5 MPLS-TP LSP protection

Linear 1-for-1 protection of MPLS-TP LSPs is supported, as defined in RFC 6378. This applies only to LSPs (not PWs).

This is supported edge-to-edge on an LSP, between two LERs, where normal traffic is transported either on the working LSP or on the protection LSP using a logical selector bridge at the source of the protected LSP.

At the sink LER of the protected LSP, the LSP that carries the normal traffic is selected, and that LSP becomes the working LSP. A protection switching coordination (PSC) protocol coordinates between the source and sink bridge, which LSP will be used, as working path and protection path. The PSC protocol is always carried on a G-ACh on the protection LSP.


The 7210 SAS supports single-phased coordination between the LSP endpoints, in which the initiating LER performs the protection switch over to the alternate path and informs the far-end LER of the switch.

Bidirectional protection switching is achieved by the PSC protocol coordinating between the two end points to determine which of the two possible paths (that is, the working or protect path), transmits user traffic at any given time.

It is possible to configure non-revertive or revertive behavior. For non-revertive, the LSP will not switch back to the working path when the PSC switch over requests end, while for revertive configurations, the LSP always returns back to the working path when the switch over requests end.


The following figures show the behavior of linear protection in more detail.

Figure 16: Normal operation



al_0233

Figure 17: Failed condition



In normal condition, user data packets are sent on the working path on both directions, from A to Z and Z to A.

A defect in the direction of transmission from node Z to node A impacts the working connection Z-to-A, and initiates the detection of a defect at the node A.

Figure 18: Failed condition - switching at A





Figure 19: Failed condition - switching at Z



The unidirectional PSC protocol initiates protection switching: the selector bridge at node A is switched to protection connection A-to-Z and the selector at node A switches to protection connection Z to-A. The PSC packet, sent from node A to node Z, requests a protection switch to node Z.

After node Z validates the priority of the protection switch request, the selector at node Z is switched to protection connection A-to-Z and the selector bridge at the node Z is switched to protection connection Z-to-A. The PSC packet, sent from node Z to node A, is used as acknowledge, informing node A about the switching.

If BFD CC or CC/CV OAM packets are used to detect defects on the working and protection path, they are inserted on both working and protection paths, and are sent regardless of whether the selected path is the currently active path.

The 7210 SAS supports the following operator commands:

- Forced Switch
- Manual Switch
- Clear
- Lockout of protection

2.3.5 Configuring MPLS-TP

This section describes the steps required to configured MPLS-TP.

2.3.5.1 Configuration overview

About this task

The following steps must be performed to configure MPLS-TP LSPs or PWs.

At the 7210 SAS LER and LSR:

Procedure

Step 1. Create an MPLS-TP context, containing nodal MPLS-TP identifiers. This is configured under `config>router>mpls>mpls-tp`.

- Step 2.** Ensure that a sufficient range of labels is reserved for static LSPs and PWs. This is configured under **config>router>mpls-labels>static-labels**.
- Step 3.** Ensure that a range of tunnel identifiers is reserved for MPLS-TP LSPs under **config>router>mpls-mpls-tp>tp-tunnel-id-range**.
- Step 4.** A user may optionally configure MPLS-TP interfaces, which are interfaces that do not use IP addressing or ARP for next hop resolution. These can only be used by MPLS-TP LSPs.

What to do next

At the 7210 SAS LER, configure:

1. OAM Templates. These contain generic parameters for MPLS-TP proactive OAM. An OAM template is configured under **config>router>mpls>mpls-tp>oam-template**.
2. BFD templates. These contain generic parameters for BFD used for MPLS-TP LSPs. A BFD template is configured under **config>router>bfd>bfd-template**.
3. Protection templates. These contain generic parameters for MPLS-TP 1-for-1 linear protection. A protection template is configured under **config>router>mpls>mpls-tp>protection-template**.
4. MPLS-TP LSPs are configured under **config>router>mpls>lsp mpls-tp**
5. Pseudowires using MPLS-TP are configured as spoke SDPs with static PW labels.

At an LSR, the user must configure an LSP transit-path under **config>router>mpls>mpls-tp>transit-path**.

The following sections describe these configuration steps in more detail.

2.3.5.2 Node-wide MPLS-TP parameter configuration

Generic MPLS-TP parameters are configured under **config>router>mpls>mpls-tp**. If a user configures **no mpls**, normally the entire mpls configuration is deleted. However, in the case of mpls-tp a check that there is no other mpls-tp configuration for example, services or tunnels using mpls-tp on the node, will be performed.

The mpls-tp context is configured as follows.

```
config
router
  mpls
    mpls-tp
      global-id <global-id>
      node-id {<ipv4address> | | <1.. .4,294,967,295>}
      [no] shutdown
    exit
```

MPLS-TP LSPs may be configured if the mpls-tp context is administratively down (shutdown), but they will remain down until the mpls-tp context is configured as administratively up. No programming of the datapath for an MPLS-TP Path occurs until the following are all true:

- MPLS-TP context is **no shutdown**
- MPLS-TP LSP context is **no shutdown**
- MPLS-TP Path context is **no shutdown**

A **shutdown** of mpls-tp will therefore bring down all MPLS-TP LSPs on the system.

The mpls-tp context cannot be deleted if MPLS-TP LSPs or SDPs exist on the system.

2.3.5.3 Node-wide MPLS-TP identifier configuration

MPLS-TP identifiers are configured for a node under the following CLI tree.

```
config
  router
    mpls
      mpls-tp
        global-id <global-id>
        node-id <node-id>
        [no] shutdown
      exit
```

The default value for the *global-id* is 0. This is used if the *global-id* is not explicitly configured. If a user expects that inter domain LSPs will be configured, then it is recommended that the global ID should be set to the local ASN of the node. as configured under **config>router**. If two-byte ASNs are used, then the most significant two bytes of the global-id are padded with zeros.

The default value of the *node-id* is the system interface IPv4 address. The MPLS-TP context cannot be administratively enabled unless at least a system interface IPv4 address is configured because MPLS requires that this value is configured.

These values are used unless overridden at the LSP or PW end-points, and apply only to static MPLS-TP LSPs and PWs.

To change the values, **config>router>mpls>mpls-tp** must be in the shutdown state. This will bring down all of the MPLS-TP LSPs on the node. New values are propagated to the system when a **no shutdown** is performed.

2.3.5.4 Static LSP and pseudowire (VC) label and tunnel ranges

SR OS reserves a range of labels for use by static LSPs and a range of labels for use by static pseudowires (SVCs); that is, LSPs and pseudowires with no dynamic signaling of the label mapping. These are configured as follows.

```
config
  router
    mpls-labels
      static-labels max-lsp-labels 991 max-svc-labels 16384
```

The minimum label value for the static LSP label starts at 32 and expands all the way to the maximum number specified. The static VC label range is contiguous with this. The dynamic label range exists above the static VC label range (the label ranges for the respective label type are contiguous). This prevents fragmentation of the label range.

The MPLS-TP tunnel ID range is configured as follows.

```
config
  router
    mpls
      mpls-tp
        [no] tp-tunnel-id-range 1 10
```

The tunnel ID range referred to here is a contiguous range of RSVP-TE Tunnel IDs is reserved for use by MPLS TP, and these IDs map to the MPLS-TP Tunnel Numbers. There are some cases where the dynamic

LSPs may have caused fragmentation to the number space such that contiguous range {max-min} is not available. In these cases, the command will fail.

There is no default value for the tunnel ID range, and it must be configured to enable MPLS-TP.

If a configuration of the tunnel ID range fails, then the system will give a reason. This could be that the initially requested range, or the change to the allocated range, is not available, that is, the tunnel IDs in that range have already been allocated by RSVP-TE. Allocated Tunnel IDs are visible using a show command.

Changing the LSP or static VC label ranges does not require a reboot.

The static label ranges for LSPs, above, apply only to static LSPs configured using the CLI tree for MPLS-TP specified in this section. Different scalability constraints apply to static LSPs configured using the following CLI introduced in earlier SAS OS releases:

config>router>mpls>static-lsp

config>router>mpls>interface>label-map

The scalability applying to labels configured using this CLI is enforced as follows:

- A maximum of 1000 static LSP names may be configured with a PUSH operation.
- A maximum of 1000 LSPs with a POP or SWAP operation may be configured.

These two limits are independent of one another, giving a combined limit of 1000 PUSH and 1000 POP/SAP operations configured on a node.

The static LSP and VC label spaces are contiguous. Therefore, the dimensioning of these label spaces requires careful planning by an operator as increasing the static LSP label space impacts the start of the static VC label space, which may already-deployed

2.3.5.5 Interface configuration for MPLS-TP

It is possible for MPLS-TP paths to use both numbered IP numbered interfaces that use ARP/static ARP, or IP unnumbered interfaces. MPLS-TP requires no changes to these interfaces. It is also possible to use a new type of interface that does not require any IP addressing or next-hop resolution.

Draft-ietf-mpls-tp-next-hop-addressing provides guidelines for the usage of various Layer 2 next-hop resolution mechanisms with MPLS-TP. If protocols such as ARP are supported, then they should be used. However, in the case where no dynamic next hop resolution protocol is used, it should be possible to configure a unicast, multicast or broadcast next-hop MAC address. The rationale is to minimize the amount of configuration required for upstream nodes when downstream interfaces are changes. A default multicast MAC address for use by MPLS-TP point-to-point LSPs has been assigned by IANA (Value: 01-00-5e-90-00-00). This value is configurable on the 7210 to support interoperability with 3rd party implementations that do not default to this value, and this no default value is implemented on the 7210.

To support these requirements, a new interface type known as an unnumbered MPLS-TP interface is introduced. This is an unnumbered interface that allows a broadcast or multicast destination MAC address to be configured. An unnumbered MPLS-TP interface is configured using the **unnumbered-mpls-tp** keyword, as follows.

```
config
router
  interface <if-name> [unnumbered-mpls-tp]
    port <port-id>[:encap-val]
    mac <local-mac-address>
    static-arp <remote-mac-addr>
    //ieee-address needs to support mcast and bcst
```

```
exit
```

The **remote-mac-address** may be any unicast, broadcast or multicast address. However, a broadcast or multicast remote-mac-address is only allowed in the **static-arp** command on Ethernet unnumbered interfaces when the **unnumbered-mpls-tp** keyword has been configured. This also allows the interface to accept packets on a broadcast or any multicast MAC address. If a packet is received with a unicast destination MAC address, then it will be checked against the configured <local-mac-address> for the interface, and dropped if it does not match. When an interface is of type **unnumbered-mpls-tp**, only MPLS-TP LSPs are allowed on that interface; other protocols are blocked from using the interface.

An unnumbered MPLS-TP interface is assumed to be point-to-point, and therefore users must ensure that the associated link is not broadcast or multicast in nature if a multicast or broadcast remote MAC address is configured.

The following is a summary of the constraints of an unnumbered MPLS-TP interface:

- It is unnumbered and may borrow/use the system interface address
- It prevents explicit configuration of a borrowed address
- It prevents IP address configuration
- It prevents all protocols except MPLS
- It prevents Deletion if an MPLS-TP LSP is bound to the Interface
- It is allowed only in network chassis mode D

MPLS-TP is only supported over Ethernet ports. The system will block the association of an MPLS-TP LSP to an interface whose port is non-Ethernet.

2.3.5.6 LER configuration for MPLS-TP

This section describes the LER configurations for MPLS-TP.

2.3.5.6.1 LSP and path configuration

MPLS-TP tunnels are configured using the **mpls-tp** LSP type at an LER under the LSP configuration, using the following CLI tree.

```
config
  router
    mpls
      lsp <xyz> [bypass-only|mpls-tp <src-tunnel-num>]
        to node-id {<a.b.c.d> | <1.. .4,294,967,295>}
        dest-global-id <global-id>
        dest-tunnel-number <tunnel-num>
        [no] working-tp-path
          lsp-num <lsp-num>
          in-label <in-label>
          out-label <out-label> out-link <if-name>
            [next-hop <ipv4-address>]
        [no] mep
          [no] oam-template <name>
          [no] bfd-enable [cc | cc_cv] // defaults to cc
          [no] shutdown
        exit
      [no] shutdown
    exit
```

```
[no] protect-tp-path
    lsp-num <lsp-num>
    in-label <in-label>
    out-label <out-label> out-link <if-name>
        [next-hop <ipv4-address> ]
[no] mep
    [no] protection-template <name>
    [no] oam-template <name>
    [no] bfd-enable [cc | cc_cv] //defaults to cc
    [no] shutdown
    exit
[no] shutdown
exit
```

<if-name> could be numbered or unnumbered interface using an Ethernet port.

<src-tunnel-num> is a mandatory create time parameter for mpls-tp tunnels, and has to be assigned by the user based on the configured range of tunnel ids. The *src-global-id* used for the LSP ID is derived from the node-wide *global-id* value configured under **config>router>mpls>mpls-tp**. A tunnel cannot be put in the **no shutdown** state unless the *global-id* is configured.

The from address of an LSP to be used in the tunnel identifier is taken to be the local node's node-id/global-id, as configured under **config>router>mpls>mpls-tp**. If that is not explicitly configured, the default value of the system interface IPv4 address is used

The **to node-id** address may be entered in 4-octet IPv4 address format or unsigned 32-bit format. This is the far-end node-id for the LSP, and does not need to be IP addresses.

The **from** and **to** addresses are used as the from and to node-id in the MPLS-TP Tunnel Identifier used for the MEP ID.

Each LSP consists of a working-tp-path and, optionally, a protect-tp-path. The protect-tp-path provides protection for the working-tp-path is 1:1 linear protection is configured (see below). Proactive OAM, such as BFD, is configured under the MEP context of each path. Protection for the LSP is configured under the protect-tp-path mep context.

The 'to' global-id is an optional parameter. If it is not entered, then the dest global ID takes the default value of 0. Global ID values of 0 are allowed and indicate that the node's configured Global ID should be used. If the local global ID value is 0, then the remote 'to' global ID must also be 0. The 'to' global ID value cannot be changed if an LSP is in use by an SDP.

The 'to' tunnel number is an optional parameter. If it is not entered, then it is taken to be the same value as the source tunnel number.

LSPs are assumed to be bidirectional and co-routed. Therefore, the system will assume that the incoming interface is the same as the out-link.

The next-hop <ip-address> can only be configured if the out-link if-name refers to a numbered IP interface. In this case, the system will determine the interface to use to reach the configured next-hop, but will check that the user-entered value for the out-link corresponds to the link returned by the system. If they do not correspond, then the path will not come up. Note that if a user changes the physical port referred to in the interface configuration, then BFD, if configured on the LSP, will go down. Users should therefore ensure that an LSP is moved to a different interface with a different port configuration to change the port that it uses. This is enforced by blocking the next-hop configuration for an unnumbered interface.

There is no check made that a valid ARP entry exists before allowing a path to come up. Therefore, a path will only be held down if BFD is down. If static ARP is not configured for the interface, then it is assumed that dynamic ARP is used. The result is that if BFD is not configured, a path can come up before ARP

resolution has completed for an interface. If BFD is not used, then it is recommended that the connectivity of the path is explicitly checked using on-demand CC/CV before sending user traffic on it.

The following is a list of additional considerations for the configuration of MPLS-TP LSPs and paths:

- The working-tp-path must be configured before the protect-tp-path.
- Likewise, the protect-tp-path has to be deleted first before the working-tp-path.
- The *lsp-num* parameter is optional. Its default value is '1' for the working-tp-path and '2' for protect-tp-path.
- The **mep** context must be deleted before a path can be deleted.
- An MPLS interface needs to be created under **config>router>mpls>interface** before using/specifying the out-label/out-link in the Forward path for an MPLS-TP LSP. Creation of the LSP will fail if the corresponding MPLS interface does not exist even though the specified router interface may be valid.
- The system will program the MPLS-TP LSP information upon a '**no shutdown**' of the TP-Path only on the very first **no shutdown**. The Working TP-Path is programmed as the Primary and the Protection TP-Path is programmed as the 'backup'.
- The system will not deprogram the IOM on an 'admin shutdown' of the MPLS-TP path. Traffic will gracefully move to the other TP-Path if valid, as determined by the proactive MPLS-TP OAM. This should not result in traffic loss. However it is recommended that the user does moves traffic to the other TP-Path through a tools command before doing 'admin shut' of an Active TP-Path.
- Deletion of the out-label/out-link sub-command under the MPLS-TP Path is not allowed after being configured. These can only be modified.
- MPLS will allow the deletion of an 'admin shutdown' TP-Path. This will cause MPLS to deprogram the corresponding TP-Path forwarding information from IOM. This can cause traffic loss for users that are bound to the MPLS-TP LSP.
- MPLS will not deprogram the IOM on a specific interface admin shut/clear unless the interface is a System Interface. However, if MPLS informs the TP-OAM module that the MPLS interface has gone down, then it triggers a switch to the standby tp-path if the associated interface went down and if it is valid.
- If a MEP is defined and shut down, then the corresponding path is also operationally down. The MEP administrative state is applicable only when a MEP is created from an MPLS-TP path.
- It is not mandatory to configure BFD or protection on an MPLS-TP path to bring the LSP up.
- If **bfd-enable cc** is configured, then CC-only mode using ACh channel 0x07 is used. If **bfd-enable cc_v** is configured, then BFD CC packets use channel 0x22 and CV packets use channel 0x23.

The protection template is associated with a LSP as a part of the MEP on the protect path. If only a working path is configured, then the protection template is not configured.

BFD cannot be enabled under the MEP context unless a named BFD template is configured.

2.3.5.6.2 Proactive CC/CV (using BFD) configuration

Generally applicable proactive OAM parameters are configured using templates.



Note:

7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T only support BFD-based CC mode. BFD-based CC-CV mode is not supported in the current release.

Proactive CC and CV uses BFD parameters such as Tx/Rx timer intervals, multiplier and other session/fault management parameters which are specific to BFD. These are configured using a BFD Template. The BFD Template may be used for non-MPLS-TP applications of BFD, and therefore contains the full set of possible configuration parameters for BFD. Only a sub-set of these may be used for any specific application.

Generic MPLS-TP OAM and fault management parameters are configured in the OAM Template.

Named templates are referenced from the MPLS-TP Path MEP configuration, so different parameter values are possible for the working and protect paths of a tunnel.

The BFD Template is configured as follows.

```
config
router
bfd
    [no] bfd-template <name>
    [no] transmit-interval <transmit-interval>
    [no] receive-interval <receive-interval>
    [no] echo-receive <echo-interval>
    [no] multiplier <multiplier>
    [no] type <cpm-np>
exit
```

The parameters are as follows:

- **transmit-interval** *transmit-interval* and the **rx** *receive-interval*: These are the transmit and receive timers for BFD packets. If the template is used for MPLS-TP, then these are the timers used by CC packets. Values are in milliseconds: 10ms to 100,000ms, with 1ms granularity. Default 10ms for CPM3 or better, 1 sec for other hardware.



Note:

For MPLS-TP CV packets, a transmit interval of 1 sec is always used.

- **multiplier** *multiplier*: Integer 3 – 20. Default: 3. This parameter is ignored for MPLS-TP combined cc-v BFD sessions, and the default of 3 used, as per RFC6428.
- **echo-receive** *echo-interval*: Sets the minimum echo receive interval, in milliseconds, for a session. Values: 100ms – 100,000ms. Default: 100. This parameter is not used by a BFD session for MPLS-TP.
- **type** *iom-hw type*: This parameter controls the system to use the IOM based hardware BFD session for the local termination point. Hardware based BFD session is enabled by default, when BFD is configured for use with MPLS-TP LSPs.

If the above BFD timer values are changed in a template, any BFD sessions on MEPs to which that template is bound will try to renegotiate their timers to the new values. BFD implementations in some MPLS-TP peer nodes may not be able handle this renegotiation, as allowed by Section 3.7.1 of RFC6428, and may take down the BFD session. This could result in unwanted behavior, for example an unexpected protection switching event. It is therefore recommended that in these circumstances, the user of 7210 SAS exercises care in modifying the BFD timer values after a BFD session is UP.

Commands within the BFD-template use a begin-commit model. To edit any value within the BFD template, a <begin> needs to be executed after the template context has been entered. However, a value will still be stored temporarily until the commit is issued. After the commit is issued, values will actually be used by other modules like the mpls-tp module and bfd module.

A BFD template is referenced from the OAM template. The OAM Template is configured as follows.

```
config
```

```
router
mpls
mpls-tp
    oam-template "state-machine-oam-template-prot"
        bfd-template "state-machine-bfd-template-prot"
        hold-time-down 0
        hold-time-up 20
    exit
```

- **hold-time-down interval:** 0-5000 deciseconds, 10ms steps, default 0. This is equivalent to the standardized hold-off timer.
- **hold-time-up interval:** 0-500 centiseconds in 100ms steps, default 2 seconds This is an additional timer that can be used to reduce BFD bouncing.
- **bfd-template name:** This is the named BFD template to use for any BFD sessions enabled under a MEP for which the OAM template is configured.

An OAM template is then applied to a MEP as described above.

2.3.5.6.3 Protection templates and linear protection configuration

Protection templates defines the generally applicable protection parameters for an MPLS-TP tunnel. Only linear protection is supported, and so the application of a named template to an MPLS-TP tunnel implies that linear protection is used.

A template is configured as follows.

```
config
router
mpls
mpls-tp
    protection-template <name>
        [no] revertive
        [no] wait-to-restore <interval>
        rapid-psc-timer <interval>
        slow-psc-timer <interval>
    exit
```

The allowed values are as follows:

- **wait-to-restore interval:** 0-720 seconds, 1 sec steps, default 300 seconds. This is applicable to revertive mode only.
- **rapid-psc-timer interval:** [10, 100, 1000ms]. Default 100ms
- **slow-psc-timer interval:** 5s-60s. Default: 5s
- **revertive:** Selects revertive behavior. Default: no revertive.

LSP Linear Protection operations are enacted using the following **tools>perform** commands.

```
tools>perform>router>mpls
    tp-tunnel
        clear {<lsp-name> | id <tunnel-id>}
        force {<lsp-name> | id <tunnel-id>}
        lockout {<lsp-name> | id <tunnel-id>}
        manual {<lsp-name> | id <tunnel-id>}
    exit
exit
```


To minimize outage times, users should use the "mpls-tp protection command" (for example, force/manual) to switch all the relevant MPLS-TP paths before executing the following commands:

- clear router mpls interface <>
- config router mpls interface <> shut

2.3.5.7 Intermediate LSR configuration for MPLS-TP LSPs

The forward and reverse directions of the MPLS-TP LSP Path at a transit LSR are configured using the following CLI tree.

```
config
  router
    mpls
      mpls-tp
        transit-path <path-name>
          [no] path-id {lsp-num <lsp-num>|working-path|protect-path

              [src-global-id <global-id>]
              src-node-id {<ipv4address> | <1.. .4,294,967,295>}
              src-tunnel-num <tunnel-num>
              [dest-global-id <global-id>]
              dest-node-id {<ipv4address> | <1.. .4,294,967,295>}
              [dest-tunnel-num <tunnel-num>]}

          forward-path
            in-label <in-label> out-label <out-label>
            out-link <if-name> [next-hop <ipv4-next-hop>]
          reverse-path
            in-label <in-label> out-label <out-label>
            [out-link <if-name> [next-hop <ipv4-next-hop>]
          [no] shutdown
```

Ensure that the *src-tunnel-num* and *dest-tunnel-num* are consistent with the source and destination of a label mapping message for a signaled LSP.

If *dest-tunnel-num* is not entered in CLI, the *dest-tunnel-num* value is taken to be the same as the *src-tunnel-num* value.

If any of the *global-id* values are not entered, the value is taken to be 0.

If the *src-global-id* value is entered, but the *dest-global-id* value is not entered, *dest-global-id* value is the same as the *src-global-id* value.

The *lsp-num* must match the value configured in the LER for a specified path. If no explicit *lsp-num* is configured, then *working-path* or *protect-path* must be specified (equating to 1 or 2 in the system).

The forward path must be configured before the reverse path. The configuration of the reverse path is optional.

The LSP-ID (path-id) parameters apply with respect to the downstream direction of the forward LSP path, and are used to populate the MIP ID for the path at this LSR.

The reverse path configuration must be deleted before the forward path.

The forward-path (and reverse-path if applicable) parameters can be configured with or without the path-id, but they must be configured if MPLS-TP OAM is to be able to identify the LSR MIP.

The transit-path can be no shutdown (as long as the forward-path/reverse-path parameters have been configured properly) with or without identifiers.

The path-id and path-name must be unique on the node. There is a one to one mapping between a specified path-name and path-id.

Traffic cannot pass through the transit-path if the transit-path is in the **shutdown** state.

2.4 RSVP

The Resource Reservation Protocol (RSVP) is a network control protocol used by a host to request specific qualities of service from the network for particular application data streams or flows. RSVP is also used by routers to deliver quality of service (QoS) requests to all nodes along the paths of the flows and to establish and maintain state to provide the requested service. RSVP requests generally result in resources reserved in each node along the datapath. MPLS leverages this RSVP mechanism to set up traffic engineered LSPs. RSVP is not enabled by default and must be explicitly enabled.

RSVP requests resources for simplex flows. It requests resources only in one direction (unidirectional). Therefore, RSVP treats a sender as logically distinct from a receiver, although the same application process may act as both a sender and a receiver at the same time. Duplex flows require two LSPs, to carry traffic in each direction.

RSVP is not a routing protocol. RSVP operates with unicast and multicast routing protocols. Routing protocols determine where packets are forwarded. RSVP consults local routing tables to relay RSVP messages.


RSVP uses two message types to set up LSPs, PATH and RESV. [Figure 20: Establishing LSPs](#) shows the process to establish an LSP.

- The sender (the ingress LER (ILER)), sends PATH messages toward the receiver, (the egress LER (ELER)) to indicate the FEC for which label bindings are wanted. PATH messages are used to signal and request label bindings required to establish the LSP from ingress to egress. Each router along the path observes the traffic type.

PATH messages facilitate the routers along the path to make the necessary bandwidth reservations and distribute the label binding to the router upstream.


- The ELER sends label binding information in the RESV messages in response to PATH messages received.
- The LSP is considered operational when the ILER receives the label binding information.

Figure 20: Establishing LSPs



The following figure shows an example of an LSP path set up using RSVP. The ingress label edge router (ILER 1) transmits an RSVP path message (path: 30.30.30.1) downstream to the egress label edge router (ELER 4). The path message contains a label request object that requests intermediate LSRs and the ELER to provide a label binding for this path.

Figure 21: LSP using RSVP path set up



In addition to the label request object, an RSVP PATH message can also contain a number of optional objects:

- Explicit route object (ERO) — When the ERO is present, the RSVP path message is forced to follow the path specified by the ERO (independent of the IGP shortest path).
- Record route object (RRO) — Allows the ILER to receive a listing of the LSRs that the LSP tunnel actually traverses.
- A session attribute object controls the path set up priority, holding priority, and local-rerouting features.

Upon receiving a path message containing a label request object, the ELER transmits a RESV message that contains a label object. The label object contains the label binding that the downstream LSR communicates to its upstream neighbor. The RESV message is sent upstream toward the ILER, in a direction opposite to that followed by the path message. Each LSR that processes the RESV message carrying a label object uses the received label for outgoing traffic associated with the specific LSP. When the RESV message arrives at the ingress LSR, the LSP is established.

2.4.1 Using RSVP for MPLS

Hosts and routers that support both MPLS and RSVP can associate labels with RSVP flows. When MPLS and RSVP are combined, the definition of a flow can be made more flexible. After an LSP is established, the traffic through the path is defined by the label applied at the ingress node of the LSP. The mapping of label to traffic can be accomplished using a variety of criteria. The set of packets that are assigned the same label value by a specific node are considered to belong to the same FEC which defines the RSVP flow.

For use with MPLS, RSVP already has the resource reservation component built-in which makes it ideal to reserve resources for LSPs.

2.4.1.1 RSVP Traffic Engineering extensions for MPLS

RSVP has been extended for MPLS to support automatic signaling of LSPs. To enhance the scalability, latency, and reliability of RSVP signaling, several extensions have been defined. Refresh messages are still transmitted but the volume of traffic, the amount of CPU utilization, and response latency are reduced while reliability is supported. None of these extensions result in backward compatibility problems with traditional RSVP implementations.

2.4.1.1.1 Hello protocol

The Hello protocol detects the loss of a neighbor node or the reset of a neighbor's RSVP state information. In standard RSVP, neighbor monitoring occurs as part of the RSVP soft-state model. The reservation state is maintained as cached information that is first installed and then periodically refreshed by the ingress and egress LSRs. If the state is not refreshed within a specified time interval, the LSR discards the state because it assumes that either the neighbor node has been lost or its RSVP state information has been reset.

The Hello protocol extension is composed of a hello message, a hello request object and a hello ACK object. Hello processing between two neighbors supports independent selection of failure detection intervals. Each neighbor can automatically issue hello request objects. Each hello request object is answered by a hello ACK object.

2.4.1.1.2 MD5 authentication of RSVP interface

When enabled on an RSVP interface, authentication of RSVP messages operates in both directions of the interface.

A node maintains a security association with its neighbors for each authentication key. The following items are stored in the context of this security association:

- The HMAC-MD5 authentication algorithm.
- Key used with the authentication algorithm.
- Lifetime of the key. A key is user-generated key using a third party software/hardware and enters the value as static string into CLI configuration of the RSVP interface. The key will continue to be valid until it is removed from that RSVP interface.
- Source Address of the sending system.
- Latest sending sequence number used with this key identifier.

The RSVP sender transmits an authenticating digest of the RSVP message, computed using the shared authentication key and a keyed-hash algorithm. The message digest is included in an Integrity object which also contains a Flags field, a Key Identifier field, and a Sequence Number field. The RSVP sender complies to the procedures for RSVP message generation in *RFC 2747, RSVP Cryptographic Authentication*.

An RSVP receiver uses the key together with the authentication algorithm to process received RSVP messages.

When a PLR node switches the path of the LSP to a bypass LSP, it does not send the Integrity object in the RSVP messages over the bypass tunnel. If an integrity object is received from the MP node, then the message is discarded because there is no security association with the next-next-hop MP node.

The 7210 SAS MD5 implementation does not support the authentication challenge procedures in RFC 2747.

2.4.2 Reservation styles

LSPs can be signaled with explicit reservation styles. A reservation style is a set of control options that specify a number of supported parameters. The style information is part of the LSP configuration. The 7210 SAS supports two reservation styles.

If FRR is enabled for the LSP and selects the facility FRR method at the head-end node, only the SE reservation style is allowed. Furthermore, if a 7210 SAS PLR node receives a path message with fast-reroute requested with facility method and the FF reservation style, it will reject the reservation. The one-to-one detour method supports both FF and SE styles.

2.4.2.1 RSVP message pacing

When a flood of signaling messages arrive because of topology changes in the network, signaling messages can be dropped which results in longer set up times for LSPs. RSVP message pacing controls the transmission rate for RSVP messages, allowing the messages to be sent in timed intervals. Pacing reduces the number of dropped messages that can occur from bursts of signaling messages in large networks.

2.4.3 RSVP overhead refresh reduction

The RSVP refresh reduction feature consists of the following capabilities implemented in accordance to *RFC 2961, RSVP Refresh Overhead Reduction Extensions*:

- **RSVP message bundling**

This capability is intended to reduce overall message handling load. The 7210 SAS supports receipt and processing of bundled message only, but no transmission of bundled messages.

- **Reliable message delivery**

This capability consists of sending a message-id and returning a message-ack for each RSVP message. It can be used to detect message loss and support reliable RSVP message delivery on a per hop basis. It also helps reduce the refresh rate because the delivery becomes more reliable.

- **Summary refresh**

This capability consists of refreshing multiples states with a single message-id list and sending negative ACKs (NACKs) for a message_id which could not be matched. The summary refresh capability reduce the amount of messaging exchanged and the corresponding message processing between peers. It does not however reduce the amount of soft state to be stored in the node.

These capabilities can be enabled on a per-RSVP-interface basis are referred to collectively as "refresh overhead reduction extensions". When the refresh-reduction is enabled on a 7210 SAS RSVP interface, the node indicates this to its peer by setting a refresh-reduction- capable bit in the flags field of the common RSVP header. If both peers of an RSVP interface set this bit, all the above three capabilities can be used. Furthermore, the node monitors the settings of this bit in received RSVP messages from the peer on the interface. As soon as this bit is cleared, the node stops sending summary refresh messages. If a peer did not set the "refresh-reduction-capable" bit, a 7210 SAS node does not attempt to send summary refresh messages.

2.4.3.1 Configuring implicit null

The implicit null label option allows a 7210 SAS egress LER to receive MPLS packets from the previous hop without the outer LSP label. The operation of the previous hop is referred to as penultimate hop popping (PHP). This option is signaled by the egress LER to the previous hop during the FEC signaling by the LDP control protocol.

The router can be configured to signal the implicit null label value over all RSVP interfaces and for all RSVP LSPs which have this node as the egress LER. In addition, the egress LER can be configured to receive MPLS packet with the implicit null label on a static LSP.

The following CLI command is used to configure the router:

config>router>ldp>implicit-null-label



Note:

RSVP must be shut down before changing the implicit null configuration option.

2.4.4 Using unnumbered Point-to-Point interface in RSVP



Note:

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

This feature introduces the use of unnumbered IP interface as a Traffic Engineering (TE) link for the signaling of RSVP P2P LSP and P2MP LSP.

An unnumbered IP interface is identified uniquely on a router in the network by the tuple {router-id, ifIndex}. Each side of the link assigns a system-wide unique interface index to the unnumbered interface. IS-IS, OSPF, RSVP, and OAM modules will use this tuple to advertise the link information, signal LSP paths over this unnumbered interface, or send and respond to an MPLS echo request message over an unnumbered interface.

The interface borrowed IP address is used exclusively as the source address for IP packets which are originated from the interface and needs to be configured to an address different from system interface for the FRR bypass LSP to come up at the ingress LER.

The borrowed IP address for an unnumbered interface is configured using the following CLI command with a default value set to the system interface address:

configure>router>interface>unnumbered [*ip-int-name* | *ip-address*].

The support of unnumbered TE link in IS-IS consists of adding a new sub-TLV of the extended IS reachability TLV, which encodes the Link Local and Link Remote Identifiers as defined in RFC 5307.

The support of unnumbered TE link in OSPF consists of adding a new sub-TLV, which encodes the same Link Local and Link Remote Identifiers in the Link TLV of the TE area opaque LSA and sends the local Identifier in the Link Local Identifier TLV in the TE link local opaque LSA as per RFC 4203.

The support of unnumbered TE link in RSVP implements the signaling of unnumbered interfaces in ERO/RRO as per RFC 3477 and the support of IF_ID RSVP_HOP object with a new Ctype as per Section 8.1.1 of RFC 3473. The IPv4 Next/Previous Hop Address field is set to the borrowed IP interface address.

The unnumbered IP is advertised by IS-IS TE and OSPF TE, and CSPF can include them in the computation of a path for a P2P LSP or for the S2L of a P2MP LSP. This feature does not, however, support defining an unnumbered interface as a hop in the path definition of an LSP.

A router creates an RSVP neighbor over an unnumbered interface using the tuple {router-id, ifIndex}. The router-id of the router which advertised a specified unnumbered interface index is obtained from the TE database. As a result, if traffic engineering is disabled in IS-IS or OSPF, a non-CSPF LSP with the next-hop for its path is over an unnumbered interface will not come up at the ingress LER because the router-id of the neighbor which has the next-hop of the path message cannot be looked up.

In this case, the LSP path will remain in operationally down state with a reason 'noRouteToDestination'. If a PATH message was received at the LSR in which traffic engineering was disabled and the next-hop for the LSP path is over an unnumbered interface, a PathErr message will be sent back to the ingress LER with the Routing Problem error code of 24 and an error value of 5 "No route available toward destination".

This feature supports all of the MPLS features available for numbered IP interfaces, with the following exceptions:

- Configuring a router-id with a value other than system.
- Signaling of an LSP path with an ERO based a loose/strict hop using an unnumbered TE link in the path hop definition.
- Signaling of one-to-one detour LSP over unnumbered interface.
- Soft preemption of LSP path using unnumbered interface.
- Inter-area LSP.
- Unnumbered RSVP interface registration with BFD.
- RSVP Hello and all Hello related capabilities such as Graceful-restart helper.
- The user SRLG database feature. The **user-srlg-db** option under MPLS allows the user to manually enter the SRLG membership of any link in the network in a local database at the ingress LER. The user cannot enter an unnumbered interface into this database and therefore all unnumbered interfaces will be considered as having no SRLG membership if the user enabled the **user-srlg-db** option.

This feature also extends the support of lsp-ping, p2mp-lsp-ping, lsp-trace, and p2mp-lsp-trace to P2P and P2MP LSPs that have unnumbered TE links in their path.

2.4.4.1 Operation of RSVP FRR facility backup over unnumbered interface

When the Point-of-Local Repair (PLR) node activates the bypass LSP by sending a PATH message to refresh the path state of protected LSP at the Merge-Point (MP) node, it must use an IPv4 tunnel sender address in the sender template object which is different than the one used by the ingress LER in the PATH message. These are the procedures specified in RFC 4090 and which are followed in the node implementation.

The node uses the address of the outgoing interface of the bypass LSP as the IPv4 tunnel sender address in the sender template object. This address will be different from the system interface address used in the sender template of the protected LSP by the ingress LER and therefore does not cause conflicts when the ingress LER acts as a PLR.



Note:

When the PLR is the ingress LER node and the outgoing interface of the bypass LSP is unnumbered, it is required that the user assigns to the interface a borrowed IP address which is different from the system interface. If not, the bypass LSP will not come up.

In addition, the PLR node will include the IPv4 RSVP_HOP object (C-Type=1) or the IF_ID RSVP_HOP object (C-Type=3) in the PATH message if the outgoing interface of the bypass LSP is numbered or unnumbered respectively.

When the MP node receives the PATH message over the bypass LSP, it will create the merge-point context for the protected LSP and associate it with the existing state if any of the following is satisfied:

- Change in C-Type of the RSVP_HOP object
- C-Type is IF_ID RSVP_HOP and did not change but IF_ID TLV is different

- Change in IPv4 Next/Previous Hop Address in RSVP_HOP object regardless of the C-Type value

These procedures at PLR and MP nodes are followed in both link-protect and node-protect FRR.

**Note:**

If the MP node is running a pre-R9.0 R4 implementation, it will reject the new IF_ID C-Type and drop the PATH over bypass. This will result in the protected LSP state expiring at the MP node, which will tear down the path. This would be the case in general when node-protect FRR is enabled and the MP node does not support unnumbered RSVP interface.

2.4.5 PCEP support for RSVP-TE LSPs

**Note:**

PCEP is only supported on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

The Path Computation Element Communication Protocol (PCEP) is one of several protocols used for communication between a wide area network (WAN) software-defined network (SDN) controller and network elements.

The 7210 SAS operates as a PCE Client (PCC) only, supporting PCC capabilities for RSVP-TE LSPs.

The following MPLS-level and LSP-level CLI commands are used to configure RSVP-TE LSPs in a router acting as a PCC.

- ```
config>router>mpls>
 pce-report rsvp-te {enable | disable}
```
- ```
config>router>mpls>lsp>
    path-profile profile-id [path-group group-id]
    pce-computation
    pce-control
    pce-report {enable | disable | inherit}
```

2.5 Traffic Engineering

Without traffic engineering, routers route traffic according to the SPF algorithm, disregarding congestion or packet types.

With traffic engineering, network traffic is routed efficiently to maximize throughput and minimize delay. Traffic engineering facilitates the mapping of traffic flows to the destination through a different (less congested) path other than the one selected by the SPF algorithm.

MPLS directs a flow of IP packets along a label switched path (LSP). LSPs are simplex, meaning that the traffic flows in one direction (unidirectional) from an ingress router to an egress router. Two LSPs are required for duplex traffic. Each LSP carries traffic in a specific direction, forwarding packets from one router to the next across the MPLS domain.

When an ingress router receives a packet, it adds an MPLS header to the packet and forwards it to the next hop in the LSP. The labeled packet is forwarded along the LSP path until it reaches the destination point. The MPLS header is removed and the packet is forwarded based on Layer 3 information such as the IP destination address. The physical path of the LSP is not constrained to the shortest path that the IGP would choose to reach the destination IP address.

2.5.1 TE metric (IS-IS and OSPF)

When the use of the TE metric is selected for an LSP, the shortest path computation after the TE constraints are applied will select an LSP path based on the TE metric instead of the IGP metric. The user configures the TE metric under the MPLS interface. Both the TE and IGP metrics are advertised by OSPF and IS-IS for each link in the network. The TE metric is part of the traffic engineering extensions of both IGP protocols.

A typical application of the TE metric is to allow CSPF to represent a dual TE topology for the purpose of computing LSP paths.

An LSP dedicated for real-time and delay sensitive user and control traffic has its path computed by CSPF using the TE metric. The user configures the TE metric to represent the delay figure, or a combined delay/jitter figure, of the link. In this case, the shortest path satisfying the constraints of the LSP path will effectively represent the shortest delay path.

An LSP dedicated for non delay sensitive user and control traffic has its path computed by CSPF using the IGP metric. The IGP metric could represent the link bandwidth or some other figure as required.

When the use of the TE metric is enabled for an LSP, CSPF will first prune all links in the network topology that do not meet the constraints specified for the LSP path. These constraints include bandwidth, admin-groups, and hop limit. CSPF will then run an SPF on the remaining links. The shortest path among the all SPF paths will be selected based on the TE metric instead of the IGP metric which is used by default. The TE metric is only used in CSPF computations for MPLS paths and not in the regular SPF computation for IP reachability.

2.5.1.1 Maintenance of TE links and nodes

Graceful shutdown is used to maintain selective links and nodes in a TE network. Before a shutdown, head-end LER nodes are notified of the imminent shutdown of the links or nodes for the purpose of maintenance, so the head-end nodes can move the paths of the LSPs away from the affected resources. Modified TE parameters for the affected links are flooded so all other nodes in the network avoid them in the CSPF calculations.

When the maintenance is over, the operator disables graceful shutdown, which reinstates and floods the user-configured TE parameters. The restored links are available for LSP path establishment.

2.5.2 Admin-group support on facility bypass backup LSP



Note:

This feature is supported only on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

This includes of the LSP primary path admin-group constraints in the computation of a fast reroute (FRR) facility bypass backup LSP so that all nodes in the LSP path protect the primary LSP path.

This feature is supported with the primary path of an RSVP P2P LSP in both intra-area and inter-area TE.

2.5.2.1 Procedures at head-end node

The user enables the signaling of the primary LSP path admin-group constraints in the FRR object at the ingress LER with the following CLI command:

config>router>mpls>lsp>fast-reroute>propagate-admin-group

When this command is enabled at the ingress LER, the admin-group constraints configured in the context of the P2P LSP primary path, or the constraints configured in the context of the LSP and inherited by the primary path, are copied into the FAST_REROUTE object. The admin-group constraints are copied into the "include-any" or "exclude-any" fields.

During LSP signaling to the downstream node, the ingress LER also propagates the admin-group constraints, which allows the node to include these constraints in the selection of the FRR backup LSP for LSP primary path protection.

The ingress LER inserts the FAST_REROUTE object, by default, in a primary LSP path message. If the user disables the object using the **config>router>mpls>no frr-object** command, the admin-group constraints are not propagated.

The same admin-group constraints can be copied into the Session Attribute object for use by LSR, typically an ABR, to expand the ERO of an inter-area LSP path. The constraints are also used by any LSR node in the path of a CSPF or non-CSPF LSP to check the admin-group constraints against the ERO, regardless of whether the hop is strict or loose. These constraints are governed strictly by the **config>router>mpls>lsp>propagate-admin-group** command.

That is, the user can copy the primary path admin-group constraints into only the FAST_REROUTE object, or only the Session Attribute object, or both.

The point-of-local repair (PLR) rules for admin-group constraint processing can use either FAST_REROUTE or Session Attribute object admin-group constraints.

2.5.2.2 Procedures at PLR node

The global **config>router>mpls>admin-group-frr** command is used to enable admin-group constraints when associating a manual or dynamic bypass LSP with the primary LSP path at a PLR node.

When the user enables this command, each PLR node reads the admin-group constraints in the FAST_REROUTE object included in the Path message of the LSP primary path. If the object is not included, the PLR reads the admin-group constraints from the Session Attribute object in the Path message.

If the PLR is also the ingress LER for the LSP primary path, the PLR uses the admin-group constraint from the LSP or path level configurations.

Whether the PLR node is also the ingress LER or just an LSR for the protected LSP primary path, the outcome of the ingress LER configuration dictates the behavior of the PLR node. The following table summarizes the bypass LSP admin-group behavior.

Table 8: Bypass LSP admin-group constraint behavior

	Ingress LER configuration	Session attribute	FRR object	Bypass LSP at PLR (LER/LSF) follows admin-group constraints
1	frr-object lsp>no propagate-admin-group	Admin color constraints not sent	Admin color constraints sent	Yes

	Ingress LER configuration	Session attribute	FRR object	Bypass LSP at PLR (LER/LSF) follows admin-group constraints
	<code>lsp>frr>propagate-admin-group</code>			
2	<code>frr-object</code> <code>lsp>propagate-admin-group</code> <code>lsp>frr>propagate-admin-group</code>	Admin color constraints sent	Admin color constraints sent	Yes
3	<code>frr-object</code> <code>lsp>propagate-admin-group</code> <code>lsp>frr>no propagate-admin-group</code>	Admin color constraints sent	Admin color constraints not sent	No
4	<code>no frr-object</code> <code>lsp>propagate-admin-group</code> <code>lsp>frr>propagate-admin-group</code>	Admin color constraints sent	Not present	Yes
5	<code>no frr-object</code> <code>lsp>no propagate-admin-group</code> <code>lsp>frr>propagate-admin-group</code>	Admin color constraints not sent	Not present	No
6	<code>no frr-object</code> <code>lsp>propagate-admin-group</code> <code>lsp>frr>no propagate-admin-group</code>	Admin color constraints sent	Not present	Yes

Next, the PLR node uses the admin-group constraints and other constraints, such as hop-limit and SRLG, to select a manual or dynamic bypass among those that are already in use.

If none of the manual or dynamic bypass LSPs satisfy the admin-group constraints and other constraints, the PLR node requests the CSPF for the path that merges closest to the protected link or node and includes or excludes the specified admin-group IDs.

Modifying the configuration of the **config>router>mpls>admin-group-frr** command does not affect existing bypass associations. The change only applies to new attempts to find a valid bypass.

2.5.3 Manual and timer resignal of RSVP-TE bypass LSP



Note:

This feature is supported only on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

The **config>router>mpls>bypass-resignal-timer** command triggers the periodic global reoptimization of all dynamic bypass LSP paths associated with an RSVP P2P LSP. The operation is performed at each expiry of the user-configurable bypass LSP resignal timer.

When this command is enabled, MPLS requests the CSPF for the best path for each dynamic bypass LSP originated on this node. The constraints, hop limit, SRLG, and admin-group constraints of the first associated LSP primary path that originally triggered the signaling of the bypass LSP must be satisfied. To do this, MPLS saves the original Path State Block (PSB) of that LSP primary path, even if the latter is torn down.

If CSPF returns no path or returns a new path with a cost that is higher than the current path, MPLS does not signal the new bypass path. If CSPF returns a new path with a cost that is lower than the current path, MPLS signals it. Also, if the new bypass path is SRLG strict disjoint with the primary path of the original PSB while the current path is SRLG loose disjoint, the manual bypass path is resigned, regardless of cost comparison.

After the new path is successfully signaled, MPLS evaluates each PSB of each PLR (that is, each unique avoid-node or avoid-link constraint) associated with the current bypass LSP path to check if the corresponding LSP primary path constraints are still satisfied by the new bypass LSP path. If so, the PSB association is moved to the new bypass LSP.

Each PSB for which the constraints are not satisfied remains associated with the PLR on the current bypass LSP and is checked at the next background PSB re-evaluation, or at the next timer or manual bypass reoptimization. Additionally, if SRLG FRR loose disjointness is configured using the **configure>router>mpls srlg-frr** command, and the current bypass LSP is SRLG disjoint with a primary path while the new bypass LSP is not SRLG disjoint, the PSB association is not moved.

If a specific PLR associated with a bypass LSP is active, the corresponding PSBs remain associated with the current PLR until the global revertive Make-Before-Break (MBB) tears down all corresponding primary paths, which also causes the current PLR to be removed.



Note:

While it is in the preceding state, the older PLR does not get any new PSB association until the PLR is removed. When the last PLR is removed, the older bypass LSP is torn down.

Additionally, PSBs that have not been moved by the dynamic or manual reoptimization of a bypass LSP as a result of the PSB constraints not being met by the new signaled bypass LSP path are re-evaluated by the FRR background task, which handles cases where the PSB has requested node protection, but its current PLR is a link-protect.

This feature is not supported with an inter-area dynamic bypass LSP and bypass LSP protecting S2L paths of a P2MP LSP.

The **tools>perform>router>mpls>resignal-bypass** command performs a manual reoptimization of a specific dynamic or manual bypass LSP, or of all dynamic bypass LSPs.

The user must configure the manual bypass LSP name. The dynamic bypass LSP name is displayed in the output of **show>router>mpls>bypass-tunnel dynamic detail**.

The **delay** option triggers the global reoptimization of all dynamic bypass LSPs at the expiry of the specified delay. This option forces the global bypass resignal timer to expire after an amount of time equal to the value of the **delay** parameter. This option has no effect on a manual bypass LSP.

However, when the bypass LSP name is specified, the named dynamic or manual bypass LSP is signaled and the associations are moved only if the new bypass LSP path has a lower cost than the current one. This behavior is different from that of the **tools>perform>router>mpls>resignal** command for the primary or secondary active path of an LSP, which signals and switches to the new path, regardless of the cost comparison. This handling is required because a bypass LSP can have a large number of PSB associations and the associated processing churn is much higher.

In the specific case where the name corresponds to a manual bypass LSP, the LSP is torn down and resignaled using the new path provided by CSPF if no PSB associations exist. If one or more PSB associations exist but no PLR is active, the **tools>perform>router>mpls>resignal-bypass** *bypass-lsp-name* configuration fails and the user is prompted to explicitly enter the **force** option. In this case, the manual bypass LSP is torn down and resignaled, temporarily leaving the associated LSP primary paths unprotected. If one or more PLRs associated with the manual bypass LSP is active, the command fails.

Finally, and as with the timer-based resignal, the PSB associations are checked for the SRLG and admin-group constraints using the updated information provided by CSPF for the current path and new path of the bypass LSP. See [RSVP-TE bypass LSP path SRLG information update in manual and timer resignal MBB](#) and [RSVP-TE bypass LSP path administrative group information update in manual and timer resignal MBB](#) for more information.

2.5.3.1 RSVP-TE bypass LSP path SRLG information update in manual and timer resignal MBB

This feature enhances procedures of the timer and manual resignal (both **delay** and **lsp** options) of the RSVP-TE bypass LSP path by updating the SRLG information of the links of the current path and checking for SRLG disjointness constraint. The following sequence describes the timer and manual resignal enhancements.

1. CSPF updates the SRLG membership of the current bypass LSP path and checks if the path violates the SRLG constraint of the first primary path that was associated with a PLR of this bypass LSP. This is referred to as the initial PSB.
2. CSPF attempts a new path computation for the bypass LSP using the initial PSB constraints.
3. MPLS uses the information returned by CSPF to determine if the new bypass path is more optimal.
 - a. If SRLG FRR strict disjointness is configured using the **configure>router>mpls> srlg-frr strict** option and CSPF indicates that the updated SRLG information of the current path violated the SRLG constraint of the PLR of the initial PSB, the new path is more optimal.
 - b. Otherwise, MPLS performs additional checks using the PLR of the initial PSB to determine if the new path is more optimal. The following table summarizes the possible cases of bypass path optimality determination.

Table 9: Determination of bypass LSP path optimality

PLR SRLG constraint check ⁹		SRLG FRR configuration (strict/loose)	Path cumulative cost comparison ⁹	More optimal path
Current path	New path			
Disjoint	Disjoint	—	New cost < current cost	New
Disjoint	Disjoint	—	New cost ≥ current cost	Current
Disjoint	Not disjoint	—	—	Current
Not disjoint	Disjoint	—	—	New
Not disjoint	Not disjoint	Strict	—	Current
Not disjoint	Not disjoint	Loose	New cost < current cost	New
Not disjoint	Not disjoint	Loose	New cost ≥ current cost	Current

4. If the path returned by CSPF is found to be a more optimal bypass path with respect to the PLR of the initial PSB, the following sequence of actions occurs.
 - a. MPLS signals and programs the new path.
 - b. MPLS moves to the new bypass path the PSB associations of all PLRs for which evaluation against the preceding table results in the new bypass path being more optimal.
 - c. Among the remaining PLRs, if the updated SRLG information of the current bypass path changed and SRLG FRR loose disjointness is configured (**configure>router>mpls>srlg-frr**), MPLS keeps this PLR PSB association with the current bypass path.
 - d. Among the remaining PLRs, if the updated SRLG information of the current bypass path changed and SRLG strict disjointness is configured (**configure>router>mpls>srlg-frr strict**), MPLS evaluates the SRLG constraint of each PLR and performs the following actions.
 - i. MPLS keeps with the current bypass path the PSB associations of all PLRs where the SRLG constraint is not violated by the updated SRLG information of the current bypass path.
These PSBs are re-evaluated at the next timer or manual resignal MBB following the procedure described in [RSVP-TE bypass LSP path SRLG information update in manual and timer resignal MBB](#).
 - ii. MPLS detaches from the current bypass path the PSB associations of all PLRs where the SRLG constraint is violated by the updated SRLG information of the current bypass path.

⁹ This check of the current path makes use of the updated SRLG and cost information provided by CSPF.

These orphaned PSBs are re-evaluated by the FRR background task, which checks unprotected PSBs on a regular basis following the procedure described in [RSVP-TE bypass LSP path SRLG information update in manual and timer resignal MBB](#).

5. If the path returned by CSPF is found to be less optimal than the current bypass path, or if CSPF does not return a new path, the following actions are performed.
 - a. If the updated SRLG information of the current bypass path did not change, MPLS keeps the current bypass path and the PSB associations of all PLRs.
 - b. If the updated SRLG information of the current bypass path changed and SRLG FRR loose disjointness is configured using the **configure>router>mpls>srlg-frr** command, MPLS keeps the current bypass path and the PSB associations of all PLRs.
 - c. If the updated SRLG information of the current bypass path has changed and SRLG strict disjointness is configured (**configure>router>mpls>srlg-frr strict**), MPLS evaluates the SRLG constraint of each PLR and performs the following actions.
 - i. MPLS keeps with the current bypass path the PSB associations of all PLRs where the SRLG constraint is not violated by the updated SRLG information of the current bypass path.

These PSBs are re-evaluated at the next timer or manual resignal MBB following the procedure described in [RSVP-TE bypass LSP path SRLG information update in manual and timer resignal MBB](#).
 - ii. MPLS detaches from the current bypass path the PSB associations of all PLRs where the SRLG constraint is violated by the updated SRLG information of the current bypass path.

These orphaned PSBs are re-evaluated by the FRR background task, which checks unprotected PSBs on a regular basis following the procedure described in [RSVP-TE bypass LSP path SRLG information update in manual and timer resignal MBB](#).

2.5.3.2 RSVP-TE bypass LSP path administrative group information update in manual and timer resignal MBB

This feature enhances procedures of the timer and manual resignal (both **delay** and **lsp** options) of a RSVP-TE bypass LSP path by updating the administrative group information of the current path links and checking for administrative group constraints. The following sequence describes the timer and manual resignal enhancements.

1. CSPF updates the administrative group membership of the current bypass LSP path and checks if the path violates the administrative group constraints of the first primary path associated with this bypass LSP. This is referred to as the initial PSB.
2. CSPF attempts a new path computation for the bypass LSP using the PLR constraints of the initial PSB.
3. MPLS uses the information returned by CSPF and determines if the new bypass path is more optimal.
 - a. If CSPF indicates the updated administrative group information of the current path violates the administrative group constraint of the initial PSB, the new path is more optimal.
 - b. Otherwise, the new path is more optimal only if its metric is lower than the updated metric of the current bypass path.
4. If the path returned by CSPF is found to be a more optimal bypass path, the following sequence of actions is performed.
 - a. MPLS signals and programs the new path.

- b. Because the administrative group constraint is not part of the PLR definition, MPLS evaluates the PSBs of all PLRs associated with the current bypass, and takes the following actions.
 - i. MPLS moves to the new bypass path the PSB associations in which the administrative group constraints are not violated by the new bypass path.
 - ii. Among the remaining PSBs, MPLS keeps with the current bypass path the PSB associations in which the administrative group constraints are not violated by the updated administrative group information of the current bypass path.

These PSBs are re-evaluated at the next timer or manual resignal MBB following the procedure described in [RSVP-TE bypass LSP path administrative group information update in manual and timer resignal MBB](#).

- iii. Among the remaining PSBs, MPLS detaches from the current bypass path the PSB associations in which the administrative group constraints are violated by the updated administrative group information of the current bypass path.

These orphaned PSBs are re-evaluated by the FRR background task, which checks unprotected PSBs on a regular basis following the procedure described in [RSVP-TE bypass LSP path administrative group information update in manual and timer resignal MBB](#).

- 5. If the path returned by CSPF is found to be less optimal than the current bypass path or if CSPF did not return a new path, the following actions are performed.
 - a. If the updated administrative group information of the current bypass path did not change, MPLS keeps the current bypass path and all PSB associations.
 - b. If the updated administrative group information of the current bypass path has changed, MPLS evaluates the PSBs of all PLRs associated with the current bypass, and performs the following actions.
 - i. MPLS keeps with the current bypass path the PSB associations in which the administrative group constraints are not violated by the updated administrative group information of the current bypass path.
 - ii. MPLS detaches from the current bypass path the PSB associations in which the administrative group constraints are violated by the updated administrative group information of the current bypass path.

These orphaned PSBs are re-evaluated by the FRR background task, which checks unprotected PSBs on a regular basis following the procedure described in [RSVP-TE bypass LSP path administrative group information update in manual and timer resignal MBB](#).

2.6 Advanced MPLS/RSVP features

This section describes the advanced MPLS and RSVP features.

2.6.1 Shared Risk Link Groups

Shared Risk Link Groups (SRLGs) is a feature that allows the user to establish a backup secondary LSP path or a FRR LSP path which is disjoint from the path of the primary LSP. Links that are members of the same SRLG represent resources sharing the same risk, for example, fiber links sharing the same conduit or multiple wavelengths sharing the same fiber.

When the SRLG option is enabled on a secondary path, CSPF includes the SRLG constraint in the computation of the secondary LSP path. This requires that the primary LSP already be established and up since the head-end LER needs the most current ERO computed by CSPF for the primary path. CSPF would return the list of SRLG groups along with the ERO during primary path CSPF computation. At a subsequent establishment of a secondary path with the SRLG constraint, the MPLS/RSVP task will query again CSPF providing the list of SLRG group numbers to be avoided. CSPF prunes all links with interfaces which belong to the same SRLGs as the interfaces included in the ERO of the primary path. If CSPF finds a path, the secondary is setup. If not, MPLS/RSVP will keep retrying the requests to CSPF.

When the SRLG option is enabled on FRR, CSPF includes the SRLG constraint in the computation of a FRR detour or bypass for protecting the primary LSP path. CSPF prunes all links with interfaces which belong to the same SRLG as the interface which is being protected, for example, the outgoing interface at the PLR the primary path is using. If one or more paths are found, the MPLS/RSVP task will select one based on best cost and will signal the bypass/detour. If not and the user included the strict option, the bypass/detour is not setup and the MPLS/RSVP task will keep retrying the request to CSPF. Otherwise, if a path exists which meets the other TE constraints, other than the SRLG one, the bypass/detour is setup.

A bypass or a detour LSP path is not guaranteed to be SRLG disjoint from the primary path. This is because only the SRLG constraint of the outgoing interface at the PLR that the primary path is using is avoided.

2.6.1.1 Enabling disjoint backup paths

The following details the steps necessary to create shared risk link groups:

- For primary/standby SRLG disjoint configuration:
 - Create an SRLG-group similar to admin groups.
 - Link the SRLG-group to MPLS interfaces.
 - Configure primary and secondary LSP paths and enable SRLG on the secondary LSP path.



Note:

The SRLG secondary LSP paths will always perform a strict CSPF query; the **srlg-frr** command is irrelevant in this case. For more information, see [srlg-frr](#).

- For FRR detours/bypass SRLG disjoint configuration:
 - Create an SRLG group, similar to admin groups.
 - Link the SRLG group to MPLS interfaces.
 - Enable the **srlg-frr** (strict/non-strict) option, which is a system-wide parameter, and it force every LSP path CSPF calculation, to take the configured SRLG memberships (and propagated through the IGP opaque-te-database) into account.
 - Configure primary FRR (one-to-one/facility) LSP paths. Consider that each PLR will create a detour/ bypass that will only avoid the SRLG memberships configured on the primary LSP path egress interface. In a one-to-one case, detour-detour merging is out of the control of the PLR, therefore the latter will not ensure that its detour will be prohibited to merge with a colliding one. For facility bypass, with the presence of several bypass type to bind to, the following priority rules will be followed:
 1. Manual bypass disjoint
 2. Manual bypass non-disjoint (eligible only if srlg-frr is non-strict)

3. Dynamic disjoint
 4. Dynamic non-disjoint (eligible only if srlg-frr is non-strict)
- Non-CSPF manual bypass is not considered.

The following figure shows a typical application of the SRLG feature is to provide for an automatic placement of secondary backup LSPs or FRR bypass/detour LSPs that minimizes the probability of fate sharing with the path of the primary LSP.

Figure 22: Shared Risk Link Groups




Fig_33

This feature is supported on OSPF and IS-IS interfaces on which RSVP is enabled.

2.6.1.2 Static configurations of SRLG memberships

This feature provides operations with the ability to enter manually the link members of SRLG groups for the entire network at any 7210 SAS node which will need to signal LSP paths (for example, a head-end node).

The operator may explicitly enable the use by CSPF of the SRLG database. In that case, CSPF will not query the TE database for IGP advertised interface SRLG information.

The SRLG secondary path computation and FRR bypass/detour path computation remains unchanged.

There are deployments where the 7210 SAS will interpret with routers that do not implement the SRLG membership advertisement through IGP SRLG TLV or sub-TLV.

In these situations, the user is provided with the ability to enter manually the link members of SRLG groups for the entire network at any 7210 SAS node which will need to signal LSP paths, for example, a head-end node.

The user enters the SRLG membership information for any link in the network by using the **interface** *ip-int-name* **srlg-group** *group-name* command in the **config>router>mpls> srlg-database>router-id** context. An interface can be associated with up to 5 SRLG groups for each execution of this command. The user can associate an interface with up to 64 SRLG groups by executing the command multiple times. The user must also use this command to enter the local interface SRLG membership into the user SRLG database. The user deletes a specific interface entry in this database by executing the **no** form of this command.

The *group-name* must have been previously defined in the SRLG **srlg-group** *group-name* **value** *group-value* command in the **config>router>mpls**. The maximum number of distinct SRLG groups the user can configure on the system is 1024.

The parameter value for *router-id* must correspond to the router ID configured under the base router instance, the base OSPF instance or the base IS-IS instance of a specific node. A single user SRLG database is maintained per node regardless if the listed interfaces participate in static routing, OSPF, IS-IS, or both routing protocols. The user can temporarily disable the use by CSPF of all interface membership information of a specific router ID by executing the **shutdown** command in the **config>router>mpls>srlg-database>router-id** context. In this case, CSPF will assume these interfaces have no SRLG membership association. The operator can delete all interface entries of a specific router ID entry in this database by executing the **no router-id** *router-address* command in the **config>router>mpls>srlg-database** context.

CSPF will not use entered SRLG membership if an interface is not listed as part of a router ID in the TE database. If an interface was not entered into the user SRLG database, it will be assumed that it does not have any SRLG membership. CSPF will not query the TE database for IGP advertised interface SRLG information.

The operator enables the use by CSPF of the user SRLG database by entering the **user-srlg-db enable** command in the **config>router>mpls** context. When the MPLS module makes a request to CSPF for the computation of an SRLG secondary path, CSPF will query the local SRLG and computes a path after pruning links which are members of the SRLG IDs of the associated primary path. Similarly, when MPLS makes a request to CSPF for a FRR bypass or detour path to associate with the primary path, CSPF queries the user SRLG database and computes a path after pruning links which are members of the SRLG IDs of the PLR outgoing interface.

The operator can disable the use of the user SRLG database by entering the **user-srlg-db disable** command in the **config>router>mpls** context. CSPF will then resumes queries into the TE database for SRLG membership information. However, the user SRLG database is maintained

The operator can delete the entire SRLG database by entering the **no srlg-database** command in the **config>router>mpls** context. In this case, CSPF will assume all interfaces have no SRLG membership association if the user has not disabled the use of this database.

2.6.2 TE graceful shutdown

Graceful shutdown provides a method to bulk re-route transit LSPs away from the node during software upgrade of a node. A solution is described in RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*. This is achieved in this draft by using a PathErr message with a specific error code Local Maintenance on TE link required flag. When a LER gets this message, it performs a make-before-break on the LSP path to move the LSP away from the links/nodes which IP addresses were indicated in the PathErr message.

Graceful shutdown can flag the affected link/node resources in the TE database so other routers will signal LSPs using the affected resources only as a last resort. This is achieved by flooding an IGP TE LSA/LSP containing link TLV for the links under graceful shutdown with the traffic engineering metric set to 0xffffffff and 0 as unreserved bandwidth.

2.6.3 Inter-area TE LSP (ERO expansion method)



Note:

This feature is only supported on the 7210 SAS-Mxp. Only P2P LSPs are supported; P2MP LSPs are not supported.

The inter-area contiguous LSP functionality provides an end-to-end TE path. Each transit node in an area can set up a TE path LSP that is based on TE information available within its local area.

A PE node that initiates an inter-area contiguous TE LSP does a partial CSPF calculation to include its local area border router (ABR) as a loose node.

An ABR that receives a Path message with loose hop ERO does a partial CSPF calculation to the next domain border router as a loose hop or CSPF to reach the final destination.


2.6.3.1 Automatic ABR node selection for inter-area LSP

The ingress LER automatically selects the ABR when setting up an inter-area RSVP P2P LSP. The user does not need to include the ABR as a loose hop in the LSP path definition.

CSPF adds the capability to compute all segments of a multi-segment intra-area or inter-area LSP path in one operation.

The following figure shows the role of each node in the signaling of an inter-area LSP with automatic ABR node selection.

Figure 23: Automatic ABR node selection for inter-area LSP




CSPF for an inter-area LSP operates as follows:

1. CSPF in the ingress LER node determines that an LSP is an inter-area LSP by doing a route lookup using the destination address of a P2P LSP (that is, the address in the to field of the LSP configuration). If there is no intra-area route to the destination address, the LSP is considered an inter-area LSP.
2. When the path of the LSP is empty, CSPF computes a single-segment intra-area path to an ABR node that advertised a prefix matching with the destination address of the LSP.
3. When the path of the LSP contains one or more hops, CSPF computes a multi-segment intra-area path that includes the hops that are in the area of the ingress LER node.
4. When all hops are in the area of the ingress LER node, the calculated path ends on an ABR node that advertised a prefix matching the destination address of the LSP.
5. When there are one or more hops that are not in the area of the ingress LER node, the calculated path ends on an ABR node that advertised a prefix matching with the first hop address that is not in the area of the ingress LER node.
6. In the case of a multi-segment inter-area LSP, if CSPF finds a hop that can be reached through an intra-area path but that resides on an ABR, CSPF calculates a path only up to that ABR. This is because there is a better chance to reach the destination of the LSP by first signaling the LSP up to that ABR and continuing the path calculation from there by having the ABR expand the remaining hops in the ERO.

The following figure shows this behavior. The TE link between ABR nodes D and E is in area 0. When node C computes the path for the LSP from C to B, in which the path specifies nodes D and E as loose hops, it fails the path computation if CSPF attempts a path all the way to the last hop in the local area, node E. Instead, CSPF stops the path at node D, which further expands the ERO by including link D to E as part of the path in area 0.

Figure 24: CSPF for an inter-area LSP



7. If there is more than one ABR that advertises a prefix, CSPF calculates a path for all ABRs. Only the shortest path is withheld. If more than one path has the shortest path, CSPF chooses a path randomly.
8. The path for an intra-area LSP cannot exit and re-enter the local area of the ingress LER.

2.6.3.1.1 Rerouting of inter-area LSP

The automatic selection of the ABR allows the ingress LER to reroute an inter-area LSP primary path through a different ABR in the following situations:

- When the local exit ABR node fails, there are two possibilities to consider:
 - The primary path is not protected at the ABR and is therefore torn down by the previous hop in the path. In this case, the ingress LER retries the LSP primary path through the ABR that has the best path for the destination prefix of the LSP.

- The primary path is protected at the ABR with a manual or dynamic bypass LSP. In this case, the ingress LER receives a Path Error message with a notification of a protection becoming active downstream and a RESV message with a Local-Protection-In-Use flag set. At the receipt of the first of these two messages, the ingress LER performs a global revertive make-before-break (MBB) to re-optimize the LSP primary path through the ABR that has the best path for the destination prefix of the LSP.
- When the local exit ABR node goes into IS-IS overload or is put into node TE graceful shutdown, the ingress LER performs an MBB to re-optimize the LSP primary path through the ABR that has the best path for the destination prefix of the LSP. The MBB is performed at the receipt of the PathErr message for the node TE shutdown or manual re-optimization of the LSP path in the case of the receipt of the IS-IS overload bit.

2.6.3.1.2 Behavior of MPLS options in inter-area LSP

Automatic ABR selection provides the following functionality:

- Path bandwidth reservation and admin-groups operate within the scope of all areas because they rely on propagating the parameter information in the Path message across the area boundary.
- The TE graceful shutdown and soft preemption functionality support MBB of the LSP path to avoid the link or node that originated the PathErr message, provided the link or node is in the local area of the ingress LER. If the PathErr originated in a remote area, the ingress LER cannot avoid the link or node when it performs the MBB because it computes the path to the local ABR exit router only. There is, however, an exception to this for the TE graceful shutdown case only: the upstream ABR nodes in the current path of the LSP record the link or node to avoid and use it in subsequent ERO expansions. This means that if the ingress LER computes a new MBB path that goes through the same exit ABR router as the current path, and all ABR upstream nodes of the node or link that originated the PathErr message are also selected in the new MBB path when the ERO is expanded, the new path avoids this link or node.
- MBB avoids the ABR node when the node is put into TE graceful shutdown.
- The **use-te-metric** option in CSPF cannot be propagated across the area boundary and operates within the scope of the local area of the ingress LER node.
- The **srlg** option on bypass LSP operates locally at each PLR within each area. The PLR node protecting the ABR checks the SRLG constraint for the path of the bypass within the local area.
- The **srlg** option on the secondary path is allowed to operate within the scope of the local area of the ingress LER node.
- The PLR node must indicate to CSPF that a request to a one-to-one detour LSP path must remain within the local area. If the destination for the detour, which is the same as that of the LSP, is outside the area, CSPF must not return a path.
- The **propagate-admin-group** option under the LSP must be enabled on the inter-area LSP if the user wants to have admin-groups propagated across the areas.
- Using automatic ABR selection, timer-based resignal of the inter-area LSP path is supported and resignals the path if the cost of the path segment to the local exit ABR changes. The cost shown for the inter-area LSP at ingress LER is the cost of the path segments to the ABR node.

2.6.3.2 Inter-area LSP support of OSPF virtual links

The OSPF virtual link extends area 0 for a router that is not connected to area 0. As a result, it makes all prefixes in area 0 reachable through an intra-area path; they are not, however, reachable because the path crosses the transit area through which the virtual link is set up to reach the area 0 remote nodes.

The TE database in a router learns all the remote TE links in area 0 from the ABR connected to the transit area, but an intra-area LSP path using these TE links cannot be signaled within area 0 because none of these links are directly connected to this node.

This inter-area LSP feature can identify when the destination of an LSP is reachable using a virtual link. In this case, CSPF automatically computes and signals an inter-area LSP through the ABR nodes that are connected to the transit area.

However, when the ingress LER for the LSP is the ABR connected to the transit area, and the destination of the LSP is the address corresponding to another ABR router ID in that same transit area, CSPF computes and signals an intra-area LSP using the transit area TE links, even when the destination router ID is only part of area 0.


2.6.3.3 Area border node FRR protection for inter-area LSP

For protection of the area border router, the upstream node of the area border router acts as a point-of-local-repair (PLR), and the next-hop node to the protected domain border routers is the merge-point (MP). Both manual and dynamic bypass are available to protect the area border node.

Manual bypass protection works only when a completely strict path is provisioned that avoids the area border node.

Dynamic bypass protection provides the automatic computation, signaling, and association with the primary path of an inter-area P2P LSP to provide ABR node protection. The following figure shows the role of each node in ABR node protection using a dynamic bypass LSP.

Figure 25: ABR protection using dynamic bypass LSP



For a PLR node within the local area of the ingress LER to provide ABR node protection, the node must dynamically signal a bypass LSP and associate it with the primary path of the inter-area LSP using the following:

- The PLR must inspect the record route object (RRO) node ID of the LSP primary path to determine the address of the node immediately downstream of the ABR in the other area.
- The PLR signals an inter-area bypass LSP with a destination address set to the address downstream of the ABR and with the exclude router object (XRO) set to exclude the node ID of the protected ABR.
- The request to CSPF is for a path to the merge-point (that is, the next-next-hop in the RRO received in the RESV for the primary path) along with the constraint to exclude the protected ABR. If CSPF returns a path that can only go to an intermediate hop, the PLR node signals the dynamic bypass and automatically includes the XRO with the address of the protected ABR. Otherwise, the PLR signals the dynamic bypass directly to the merge-point node with no XRO object in the Path message.
- If a node-protect dynamic bypass cannot be found or signaled, the PLR node attempts a link-protect dynamic bypass LSP. As in existing implementation of dynamic bypass within the same area, the PLR attempts in the background to signal a node-protect bypass at the receipt of every third RESV refresh message for the primary path.
- Refresh reduction over dynamic bypass works only if the RRO node ID also contains the interface address. Otherwise, the neighbor is not created when the bypass is activated by the PLR. The Path state times out after three refreshes following the activation of the bypass backup LSP.

A one-to-one detour backup LSP cannot be used at the PLR for the protection of the ABR. As a result, a PLR node does not signal a one-to-one detour LSP for ABR protection. In addition, an ABR rejects a Path message that is received from a third party implementation with a detour object and with the ERO having the next-hop loose. This is performed regardless of whether the **cspf-on-loose-hop** command is enabled on the node. That is, the router as a transit ABR for the detour path rejects the signaling of an inter-area detour backup LSP.

2.7 Point-to-Multipoint (P2MP) LSP



Note:

- This feature is supported on all 7210 SAS platforms as described in this document, except the 7210 SAS-Sx 1/10GE and 7210 SAS-Sx 10/100GE, and platforms operating in access-uplink mode.
- P2MP LSPs signaled using RSVP or mLDP is only supported on 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.
- Inter-area RSVP-TE P2MP LSPs are not supported.


The following is a generic description of the P2MP LSPs functionality.

Point-to-multipoint (P2MP) LSP allows the source of multicast traffic to forward packets to one or many multicast receivers over a network without requiring a multicast protocol, such as PIM, to be configured in the network core routers. A P2MP LSP tree is established in the control plane which path consists of a head-end node, one or many branch nodes, and the leaf nodes. Packets injected by the head-end node are replicated in the data plane at the branching nodes before they are delivered to the leaf nodes.

2.7.1 Application in video broadcast

The following figure shows the use of the SR product family in triple play application (TPSDA). 7210 SAS devices could be attached to the BSA device as part of the access network.

Figure 26: Application of P2MP LSP in video broadcast



A PIM-free core network can be achieved by deploying P2MP LSPs using other core routers. The router can act as the ingress LER receiving the multicast packets from the multicast source and forwarding them over the P2MP LSP. A router can act as a leaf for the P2MP LSP tree initiated from the head-end router co-located with the video source. The router can also act as a branch node serving other leaf nodes and supports the replication of multicast packets over P2MP LSPs.

2.7.2 P2MP LSP data plane

A P2MP LSP is a unidirectional label switched path (LSP) which inserts packets at the root (ingress LER) and forwards the exact same replication of the packet to one or more leaf nodes (egress LER). The packet can be replicated at the root of P2MP LSP tree and, or at a transit LSR which acts as a branch node for the P2MP LSP tree.

The data link layer code-point, for example Ethertype when Ethernet is the network port, continues to use the unicast codepoint defined in RFC 3032, MPLS Label Stack Encoding, and which is used on P2P LSP. This change is specified in draft-ietf-mpls-multicast-encaps, MPLS Multicast Encapsulations.

2.7.2.1 Procedures at ingress LER node

The following procedures occur at the root of the P2MP LSP (head-end or ingress LER node):

- First, the P2MP LSP state is established via the control plane. Each leaf of the P2MP LSP will have a next-hop label forwarding entry (NHLFE) configured in the forwarding plane for each outgoing interface.

- The user maps a specific multicast destination group address to the P2MP LSP in the base router instance by configuring a static multicast group under a tunnel interface representing the P2MP LSP.
- An FTN entry is programmed at the ingress of the head-end node that maps the FEC of a received user IP multicast packet to a list of outgoing interfaces (OIF) and corresponding NHLFEs.
- The head-end node replicates the received IP multicast packet to each NHLFE. Replication is performed at ingress toward the fabric or at egress forwarding engine depending on the location of the OIF.
- At ingress, the head-end node performs a PUSH operation on each of the replicated packets.

2.7.2.2 Procedures at LSR node

The following procedures occur at an LSR node that is not a branch node:

- The LSR performs a label swapping operation on a leaf of the P2MP LSP. This is a conventional operation of an LSR in a P2P LSP. An ILM entry is programmed at the ingress of the LSR to map an incoming label to a NHLFE. The following is an exception handling procedure for control packets received on an ILM in an LSR.
- Packets that arrive with the TTL in the outer label expiring are sent to the CPM for further processing and are not forwarded to the egress NHLFE.

2.7.2.3 Procedures at branch LSR node

The following procedures occur at an LSR node that is a branch node:

- The LSR performs a replication and a label swapping for each leaf of the P2MP LSP. An ILM entry is programmed at the ingress of the LSR to map an incoming label to a list of OIF and corresponding NHLFEs.
- There is a system defined limit on the number of OIF/NHLFEs per ILM entry.

For control packets received on an ILM in a branch LSR, packets that arrive with the TTL in the outer label expiring are sent to the CPM for further processing and not copied to the LSP branches.

2.7.2.4 Procedures at egress LER node

At the leaf node of the P2MP LSP (egress LER), the egress LER performs a pop operation. An ILM entry is programmed at the ingress of the egress LER to map an incoming label to a list of next-hop/OIF.

For control packets received on an ILM in an egress LER, the packet is sent to the CPM for further processing if there is any of the IP header exception handling conditions set after the label is popped: 127/8 destination address, router alert option set, or any other options set.

2.7.2.5 Procedures at BUD LSR node

At an LSR node which is both a branch node and an egress leaf node (bud node), the bud LSR performs a pop operation on one or many replications of the received packet and a swap operation of the remaining replications. An ILM entry is programmed at ingress of the LSR to map the incoming label to list of NHLFE/OIF and next-hop/OIF. The exact same packets are replicated to an LSP leaf and to a local interface.

The following are the exception handling procedures for control packets received on an ILM in a bud LSR:

- Packets which arrive with the TTL in the outer label expiring are sent to the CPM and are not copied to the LSP branches.
- Packets whose TTL does not expire are copied to all branches of the LSP. The local copy of the packet is sent to the CPM for further processing if there is any of the IP header exception handling conditions set after the label is popped: 127/8 destination address, router alert option set, or any other options set.

2.7.3 RSVP control plane in a P2MP LSP

P2MP RSVP LSP is specified in RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*.

A P2MP LSP is modeled as a set of root-to-leaf (S2L) sub-LSPs. The root, for example the head-end node, triggers signaling using one or multiple path messages. A path message can contain the signaling information for one or more S2L sub-LSPs. The leaf sub-LSP paths are merged at branching points.

A P2MP LSP is identified by the combination of <P2MP ID, tunnel ID, extended tunnel ID> part of the P2MP session object, and <tunnel sender address, LSP ID> fields in the P2MP sender_template object.

A specific sub-LSP is identified by the <S2L sub-LSP destination address> part of the S2L_SUB_LSP object and an ERO and secondary ERO (SERO) objects.

The following are characteristics of this feature:

- Supports the deaggregated method for signaling the P2MP RSVP LSP. Each root to leaf is modeled as a P2P LSP in the RSVP control plane. Only data plane merges the paths of the packets.
- Each S2L sub-LSP is signaled in a separate path message. Each leaf node responds with its own resv message. A branch LSR node will forward the path message of each S2L sub-LSP to the downstream LSR without replicating it. It will also forward the resv message of each S2L sub-LSP to the upstream LSR without merging it with the resv messages of other S2L sub-LSPs of the same P2MP LSP. The same is done for subsequent refreshes of the path and resv states.
- The node will drop aggregated RSVP messages on the receive side if originated by another vendor's implementation.
- The user configures a P2MP LSP by specifying the optional create-time parameter **p2mp-lsp** following the LSP name. Next, the user creates a primary P2MP instance using the keyword **primary-p2mp-instance**. Then a path name of each S2L sub-LSP must be added to the P2MP instance using the keyword **s2l-path**. The paths can be empty paths or can specify a list of explicit hops. The path name must exist and must have been defined in the **config>router>mpls>path** context.
- The same path name can be re-used by more than one S2L of the primary P2MP instance. However the keyword must have a unique argument per S2L as it corresponds to the address of the egress LER node.
- The user can configure a secondary instance of the P2MP LSP to backup the primary one. In this case, the user enters the name of the secondary P2MP LSP instance under the same LSP name. One or more secondary instances can be created. The trigger for the head-end node to switch the path of the LSP from the primary P2MP instance to the secondary P2MP instance is to be determined. This could be based on the number of leaf LSPs which went down at a specified time.
- The following parameters can be used with a P2MP LSP: **adaptive**, **cspf**, **exclude**, **fast-reroute**, **from**, **hop-limit**, **include**, **metric**, **retry-limit**, **retry-timer**, **resignal-timer**.
- The following parameters cannot be used with a P2MP LSP: **adspec**, **primary**, **secondary**, **to**.

- The **to** parameter is not available at the LSP level but at the path level of each S2L sub-LSP of the primary or secondary instance of this P2MP LSP.
- The node ingress LER will not inset an adspec object in the path message of an S2L sub-LSP. If received in the resv message, it will be dropped. The operational MTU of an S2L path is derived from the MTU of the outgoing interface of that S2L path.
- The hold-timer configured in the **config>router>mpls>hold-timer** context applies when signaling or re-signaling an individual S2L sub-LSP path. It does not apply when the entire tree is signaled or re-signaled.
- The head-end node can add or remove a S2L sub-LSP of a specific leaf node without impacting forwarding over the already established S2L sub-LSPs of this P2MP LSP and without re-signaling them.
- The head-end node performs a make-before break (MBB) on an individual S2L path of a primary P2MP instance whenever it applies the FRR global revertive procedures to this path. If CSPF finds a new path, RSVP signals this S2L path with the same LSP-ID as the existing path.
- All other configuration changes, such as adaptive/no-adaptive (when an MBB is in progress), use-te-metric, no-frr, cspf/no-cspf, result in the tear-down and re-try of all affected S2L paths.
- MPLS requests CSPF to re-compute the whole set of S2L paths of a specified active P2MP instance each time the P2MP re-signal timer expires. The P2MP re-signal timer is configured separately from the P2P LSP. MPLS performs a global MBB and moves each S2L sub-LSP in the instance into its new path using a new P2MP LSP ID if the global MBB is successful. This is regardless of the cost of the new S2L path.
- MPLS will request CSPF to re-compute the whole set of S2L paths of a specified active P2MP instance each time the user performs a manual re-signal of the P2MP instance. MPLS then always performs a global MBB and moves each S2L sub-LSP in the instance into its new path using a new P2MP LSP ID if the global MBB is successful. This is regardless of the cost of the new S2L path. The user executes a manual re-signal of the P2MP LSP instance using the command:
tools>perform>router>mpls>resignal p2mp-lsp *lsp-name* p2mp-instance *instance-name*.
- When performing global MBB, MPLS runs a separate MBB on each S2L in the P2MP LSP instance. If an S2L MBB does not succeed the first time, MPLS will re-try the S2L using the re-try timer and re-try count values inherited from P2MP LSP configuration. However, there will be a global MBB timer set to 600 seconds and which is not configurable. If the global MBB succeeds, for example, all S2L MBBs have succeeded, before the global timer expires, MPLS moves the all S2L sub-LSPs into their new path. Otherwise when this timer expires, MPLS checks if all S2L paths have at least tried once. If so, it then aborts the global MBB. If not, it will continue until all S2Ls have re-tried once and then aborts the global MBB. Once global MBB is aborted, MPLS will move all S2L sub-LSPs into the new paths only if the set of S2Ls with a new path found is a superset of the S2Ls which have a current path which is up.
- While make-before break is being performed on individual S2L sub-LSP paths, the P2MP LSP will continue forwarding packets on S2L sub-LSP paths which are not being re-optimized and on the older S2L sub-LSP paths for which make-before-break operation was not successful. MBB will therefore result in duplication of packets until the old path is torn down.
- The MPLS datapath of an LSR node, branch LSR node, and bud LSR node will be able to re-merge S2L sub-LSP paths of the same P2MP LSP in case their ILM is on different incoming interfaces and their NHLFE is on the same or different outgoing interfaces. This could occur anytime there are equal cost paths through this node for the S2L sub-LSPs of this P2MP LSP.
- Link-protect FRR bypass using P2P LSPs is supported. In link protect, the PLR protecting an interface to a branch LSR will only make use of a single P2P bypass LSP to protect all S2L sub-LSPs traversing the protected interface.

- Refresh reduction on RSVP interface and on P2P bypass LSP protecting one or more S2L sub-LSPs.
- A manual bypass LSP cannot be used for protecting S2L paths of a P2MP LSP.
- The following MPLS features do operate with P2MP LSP:
 - BFD on RSVP interface
 - MD5 on RSVP interface
 - IGP metric and TE metric for computing the path of the P2MP LSP with CSPF
 - SRLG constraint for computing the path of the P2MP LSP with CSPF. SRLG is supported on FRR backup path only.
 - TE graceful shutdown
 - Admin group constraint
- The following MPLS features are not operable with P2MP LSP:
 - Class based forwarding over P2MP RSVP LSP
 - LDP-over-RSVP where the RSVP LSP is a P2MP LSP
 - Diff-Serv TE
 - Soft preemption of RSVP P2MP LSP

2.7.4 Forwarding multicast packets over RSVP P2MP LSP in the base router

Multicast packets are forwarded over the P2MP LSP at the ingress LER based on a static join configuration of the multicast group against the tunnel interface associated with the originating P2MP LSP. At the egress LER, packets of a multicast group are received from the P2MP LSP via a static assignment of the specific <S,G> to the tunnel interface associated with a terminating LSP.

2.7.4.1 Procedures at ingress LER node

To forward multicast packets over a P2MP LSP, perform the following steps:

1. Create a tunnel interface associated with the P2MP LSP: **config>router>tunnel-interface rsvp-p2mp lsp-name**. (The **config>router>pim>tunnel-interface** command has been discontinued.)
2. Add static multicast group joins to the PIM interface, either as a specific <S,G> or as a <*,G>:
config>router>igmp>tunnel-if>static>group>source ip-address and **config>router>igmp>tunnel-if>static>group>starg**.

The tunnel interface identifier consists of a string of characters representing the LSP name for the RSVP P2MP LSP. MPLS will pass a more structured tunnel interface identifier to PIM. The structure will follow the one BGP uses to distribute the PMSI tunnel information in BGP multicast VPN as specified in draft-ietf-l3vpn-2547bis-mcast-bgp, *Multicast in MPLS/BGP IP VPNs*. The format is: <extended tunnel ID, reserved, tunnel ID, P2MP ID> as encoded in the RSVP-TE P2MP LSP session_attribute object in RFC 4875.

The user can create one or more tunnel interfaces in PIM and associate each to a different RSVP P2MP LSP. The user can then assign static multicast group joins to each tunnel interface. A specific <*,G> or <S,G> can only be associated with a single tunnel interface.

A multicast packet which is received on an interface and which succeeds the RPF check for the source address will be replicated and forwarded to all OIFs which correspond to the branches of the P2MP LSP. The packet is sent on each OIF with the label stack indicated in the NHLFE of this OIF. The packets will

also be replicated and forwarded natively on all OIFs which have received IGMP or PIM joins for this <S,G>.

The multicast packet can be received over a PIM or IGMP interface which can be an IES interface, a spoke SDP-terminated IES interface, or a network interface.

To duplicate a packet for a multicast group over the OIF of both P2MP LSP branches and the regular PIM or IGMP interfaces, the tap mask for the P2MP LSP and that of the PIM based interfaces will need to be combined into a superset MCID.

2.7.4.2 Procedures with a primary tunnel at egress LER node

The user configures a tunnel interface and associates it with a terminating P2MP LSP leaf using the command: **config>router>tunnel-interface rsvp-p2mp lsp-name sender sender-address**. (The **config>router>pim>tunnel-interface** command has been discontinued).

The tunnel interface identifier consists of strings of characters representing the LSP name for the RSVP P2MP LSP followed by the system address of the ingress LER. The LSP name must correspond to a P2MP LSP name configured by the user at the ingress LER and must not contain the special character ":". MPLS will pass a more structured tunnel interface identifier to PIM. The structure will follow the one BGP uses to distribute the PMSI tunnel information in BGP multicast VPN as specified in draft-ietf-l3vpn-2547bis-mcast-bgp. The format is: <extended tunnel ID, reserved, tunnel ID, P2MP ID> as encoded in the RSVP-TE P2MP LSP session_attribute object in RFC 4875.

The egress LER accepts multicast packets using the following methods:

1. The regular RPF check on unlabeled IP multicast packets, which is based on routing table lookup.
2. The static assignment which specifies the receiving of a multicast group <*,G> or a specific <S,G> from a primary tunnel-interface associated with an RSVP P2MP LSP.

One or more primary tunnel interfaces in the base router instance can be configured. In other words, the user will be able to receive different multicast groups, <*,G> or specific <S,G>, from different P2MP LSPs. This assumes that the user configured static joins for the same multicast groups at the ingress LER to forward over a tunnel interface associated with the same P2MP LSP.

A multicast info policy CLI option allows the user to define a bundle and specify channels in the bundle that must be received from the primary tunnel interface. The user can apply the defined multicast info policy to the base router instance.

At any specified time, packets of the same multicast group can be accepted from either the primary tunnel interface associated with a P2MP LSP or from a PIM interface. These are mutually exclusive options. As soon as a multicast group is configured against a primary tunnel interface in the multicast info policy, it is blocked from other PIM interfaces.

However, if the user configured a multicast group to be received from a specified primary tunnel interface, there is nothing preventing packets of the same multicast group from being received and accepted from another primary tunnel interface. However, an ingress LER will not allow the same multicast group to be forwarded over two different P2MP LSPs. The only possible case is that of two ingress LERs forwarding the same multicast group over two P2MP LSPs toward the same egress LER.

A multicast packet received on a tunnel interface associated with a P2MP LSP can be forwarded over a PIM or IGMP interface which can be an IES interface, a spoke SDP-terminated IES interface, or a network interface.

Packets received from a primary tunnel-interface associated with a terminating P2MP LSP cannot be forwarded over a tunnel interface associated with an originating P2MP LSP.


2.7.5 Configuration guidelines for RSVP P2MP LSPs

- Before using P2MP LSPs with NG-MVPN, resources must be allocated from the sf-ingress-internal-tcam resource pool using the **configure>system>global-res-profile>sf-ingress-internal-tcam>mpls-p2mp** command. In addition, if the 7210 SAS-R6 is deployed as a bud router, the **configure>system>loopback-no-svc-port p2mpbud p2mpbud-port-id** command must be used to configure one of the front-panel ports as a loopback port.
- Ingress FC classification is available for packets received on a P2MP LSP on a network port IP interface that needs to be replicated to IP receivers. Ingress FC classification allows users to prioritize multicast traffic to IP receivers in the service. Also available is the capability to mark the packet with IP DSCP values while sending the multicast stream out of the IP interface. To enable ingress FC classification, use the **loopback-no-svc-port [p2mpbud p2mpbud-port-id [classification]]** command. Before using the command, users must ensure that sufficient resources are available in the network port ingress CAM resource pool and MPLS EXP ingress profile map resource pool. The **tools>dump>system-resources** command can be used to check resource availability.

2.8 MPLS/RSVP configuration process overview

The following figure shows the process to configure MPLS and RSVP parameters.

Figure 27: MPLS and RSVP configuration and implementation flow



2.9 Configuration notes

This section describes MPLS and RSVP restrictions.

- Interfaces must already be configured in the **config>router>interface** context before they can be specified in MPLS and RSVP.
- A router interface must be specified in the **config>router>mpls** context to apply it or modify parameters in the **config>router>rsvp** context.
- A system interface must be configured and specified in the **config>router>mpls** context.
- Paths must be created before they can be applied to an LSP.

2.10 Configuring MPLS and RSVP with CLI

This section provides information to configure MPLS and RSVP using the command line interface.

2.11 MPLS configuration overview

Multiprotocol Label Switching (MPLS) enables routers to forward traffic based on a simple label embedded into the packet header. A router examines the label to determine the next hop for the packet, saving time for router address lookups to the next node when forwarding packets. MPLS is not enabled by default and must be explicitly enabled.

2.11.1 LSPs

To configure MPLS-signaled label switched paths (LSPs), an LSP must run from an ingress router to an egress router. Configure only the ingress router and configure LSPs to allow the software to make the forwarding decisions or statically configure some or all routers in the path. The LSP is set up by Resource Reservation Protocol (RSVP), through RSVP signaling messages. The 7210 SAS automatically manages label values. Labels that are automatically assigned have values ranging from 1,024 through 1,048,575 (see [Label values](#)).

A static LSP is a manually set up LSP where the next-hop IP address and the outgoing label are explicitly specified.

2.11.2 Paths

To configure signaled LSPs, you must first create one or more named paths on the ingress router. For each path, the transit routers (hops) in the path are specified.

2.11.3 Router interface

At least one router interface and one system interface must be defined in the **config>router>interface** context to configure MPLS on an interface.

2.11.4 Choosing the signaling protocol

If only static label switched paths are used in your configurations, then you must manually define the paths through the MPLS network. Label mappings and actions configured at each hop must be specified. You do not need to enable RSVP if you are configuring static LSPs.

If dynamic LSP signaling is implemented in your network, then RSVP must be specified. Enable signaling protocols only on the links where the functionality is required.

To implement MPLS, the following entities must be enabled:

- MPLS must be enabled on all routers that are part of an LSP.
- RSVP must be enabled on the same routers.

When MPLS is enabled and either RSVP is also enabled, MPLS uses RSVP to set up the configured LSPs. For example, when you configure an LSP with both MPLS and RSVP running, RSVP initiates a session for the LSP. RSVP uses the local router as the RSVP session sender and the LSP destination as the RSVP session receiver. When the RSVP session is created, the LSP is set up on the path created by the session. If the session is not successfully created, RSVP notifies MPLS, MPLS can then either initiate backup paths or retry the initial path.

2.12 Basic MPLS configuration

This section provides information to configure MPLS and provides configuration examples of common configuration tasks. To enable MPLS on the 7210 SAS, you must configure at least one MPLS interface. The other MPLS configuration parameters are optional. The following is an example of an MPLS configuration.

The **admin-group** is configured in the **config>router>if-attribute** context and associated with the MPLS interface in the **config>router>mpls>interface** context.

Example

```
ALA-1>config>router>if-attr# info
-----
admin-group "green" value 15
admin-group "yellow" value 20
admin-group "red" value 25
-----

A:ALA-1>config>router>mpls# info
#-----
echo "IP Configuration"
#-----
if-attribute
    admin-group "green" 15
    admin-group "yellow" 20
    admin-group "red" 25
    interface "system"
    exit
    interface "StaticLabelPop"
        admin-group "green"
        label-map 50
        pop
        no shutdown
```



```
        exit
    exit
    interface "StaticLabelPop"
        label-map 35
            swap 36 nexthop 10.10.10.91
            no shutdown
        exit
    exit
    path "secondary-path"
        no shutdown
    exit
    path "to-NYC"
        hop 1 10.10.10.104 strict
        no shutdown
    exit
    lsp "lsp-to-eastcoast"
        to 10.10.10.104
        from 10.10.10.103
        fast-reroute one-to-one
        exit
        primary "to-NYC"
        exit
        secondary "secondary-path"
        exit
        no shutdown
    exit
    static-lsp "StaticLabelPush"
        to 10.10.11.105
        push 60 nexthop 10.10.11.105
        no shutdown
    exit
    no shutdown
-----
A:ALA-1>config>router>mpls#
```

2.13 Common configuration tasks

This section provides a brief overview of the tasks to configure MPLS and provides the CLI commands.

The following protocols must be enabled on each participating router:

- MPLS
- RSVP (for RSVP-signaled MPLS only)
- LDP

In order for MPLS to run, you must configure at least one MPLS interface in the **config>router>mpls** context.

- An interface must be created in the **config>router>interface** context before it can be applied to MPLS.
- In the **config>router>mpls** context, configure path parameters for configuring LSP parameters. A path specifies some or all hops from ingress to egress. A path can be used by multiple LSPs.
- When an LSP is created, the egress router must be specified in the **to** command and at least one primary or secondary path must be specified. All other statements under the LSP hierarchy are optional.

2.13.1 Configuring global MPLS parameters

Admin groups signify link colors, such as red, yellow, or green. MPLS interfaces advertise the link colors that they support. CSPF uses the information when paths are computed for constraint-based LSPs. CSPF must be enabled in order for admin groups to be relevant.

Use the following syntax to configure MPLS admin-group parameters.

```
Config>router>if-attribute
    admin-group group-name value group-value
mpls
    frr-object
    resignal-timer minutes
```

Example: Admin group configuration output

```
ALA-1>config>router>if-attr# info
-----
admin-group "green" value 15
admin-group "yellow" value 20
admin-group "red" value 25
-----
A:ALA-1>config>router>mpls# info
-----
                resignal-timer 500
...
-----
A:ALA-1>config>router>mpls#
```

2.13.2 Configuring an MPLS interface

Configure the **label-map** parameters if the interface is used in a static LSP. Use the following syntax to configure an MPLS interface on a router.

```
config>router>mpls
interface
    no shutdown
    admin-group group-name [group-name...(up to 32 max)]
    label-map
        pop
        swap
        no shutdown
    srlg-group group-name [group-name...(up to 5 max)]
    te-metric value
```

Example: Interface configuration output

```
A:ALA-1>config>router>mpls# info
-----
...
        interface "to-104"
            admin-group "green"
            admin-group "red"
            admin-group "yellow"
            label-map 35
            swap 36 nexthop 10.10.10.91
```

```
        no shutdown
      exit
    exit
    no shutdown
  ...
-----
A:ALA-1>config>router>mpls#
```

2.13.3 Configuring MPLS paths

Configure an LSP path to use in MPLS. When configuring an LSP, the IP address of the hops that the LSP should traverse on its way to the egress router must be specified. The intermediate hops must be configured as either **strict** or **loose** meaning that the LSP must take either a direct path from the previous hop router to this router (**strict**) or can traverse through other routers (**loose**).

Use the following syntax to configure a path.

```
config>router> mpls
path path-name
hop hop-index ip-address {strict|loose}
no shutdown
```

Example: Path configuration output

```
A:ALA-1>config>router>mpls# info
-----
      interface "system"
      exit
      path "secondary-path"
        hop 1 10.10.0.121 strict
        hop 2 10.10.0.145 strict
        hop 3 10.10.0.1 strict
        no shutdown
      exit
      path "to-NYC"
        hop 1 10.10.10.103 strict
        hop 2 10.10.0.210 strict
        hop 3 10.10.0.215 loose
      exit
-----
A:ALA-1>config>router>mpls#
```

2.13.4 Configuring an MPLS LSP

Configure an LSP path for MPLS. When configuring an LSP, you must specify the IP address of the egress router in the **to** statement. Specify the primary path to be used. Secondary paths can be explicitly configured or signaled upon the failure of the primary path. All other statements are optional.

Example: MPLS LSP configuration output

```
A:ALA-1>config>router>mpls# info
-----
...
      lsp "lsp-to-eastcoast"
        to 192.168.200.41
        rsvp-resv-style ff
```

```
        cspf
        include "red"
        exclude "green"
        adspec
        fast-reroute one-to-one
        exit
        primary "to-NYC"
            hop-limit 10
        exit
        secondary "secondary-path"
            bandwidth 50000
        exit
        no shutdown
    exit
    no shutdown
-----
A:ALA-1>config>router>mpls#
```

2.13.4.1 Configuring a static LSP

An LSP can be explicitly (statically) configured. Static LSPs are configured on every node along the path. The label's forwarding information includes the address of the next hop router.

Use the following syntax to configure a static LSP.

```
config>router>mpls
static-lsp lsp-name
to ip-address
push out-label nexthop ip-addr
no shutdown
```

Example: Static LSP configuration output


```
A:ALA-1>config>router>mpls# info
-----
...
        static-lsp "static-LSP"
            to 10.10.10.124
            push 60 nexthop 10.10.42.3
            no shutdown
        exit
...
-----
A:ALA-1>config>router>mpls#
```

2.13.5 Configuring manual bypass tunnels

About this task

Consider the network setup in the following figure that shows nodes A through F.

Figure 28: Manual bypass tunnels



Procedure

Step 1. The user first configures the option to disable the dynamic bypass tunnels.

Listed below are the steps to configure the manual bypass tunnels:

- a. Configure the option to disable the dynamic bypass tunnels on the 7210 SAS node B (if required). The CLI for this configuration is: **config>router>mpls>dynamic-bypass [disable | enable]** The dynamic bypass tunnels are enabled by default.
- b. Configure an LSP on node B, such as B-E-F-C which is used only as bypass. The user specifies each hop in the path, for example, the bypass LSP has a strict path.

Including the **bypass-only** keyword disables the following options under the LSP configuration:

- bandwidth
- fast-reroute
- secondary

The following LSP configuration options are allowed:

- adaptive
- adspec
- cspf
- exclude
- hop-limit
- include
- metric

Example

Bypass tunnel configuration output

```

A:7210 SAS>config>router>mpls>path# info
-----
...
path "BEFC"
    hop 10 10.10.10.11 strict
    hop 20 10.10.10.12 strict
    hop 30 10.10.10.13 strict
    no shutdown
exit

lsp "bypass-BC"
    to 10.10.10.15
    primary "BEFC"
    exit
    no shutdown
...
    
```

```
-----  
A:7210 SAS >config>router>mpls>path#
```

Step 2. Configure an LSP from A to D and indicate fast-reroute bypass protection, select the facility as "FRR method" (**config>router>mpls>lsp>fast-reroute facility**).

Observe if the following criteria apply:

- If the LSP passes through B
- A bypass is requested
- The next hop is C
- A manually configured bypass-only tunnel exists from B to C (excluding link B to C)

Expected outcome

Node B uses the manually configured bypass-only tunnel from B to C.

2.14 Configuring RSVP parameters

RSVP is used to set up LSPs. RSVP must be enabled on the router interfaces that are participating in signaled LSPs. The **keep-multiplier** and **refresh-time** default values can be modified in the RSVP context.

Initially, interfaces are configured in the **config>router>mpls>interface** context. Only these existing (MPLS) interfaces are available to modify in the **config>router>rsvp** context. Interfaces cannot be directly added in the RSVP context.

Example: RSVP configuration output

```
A:ALA-1>config>router>rsvp# info  
-----  
interface "system"  
    no shutdown  
    exit  
    interface to-104  
        hello-interval 4000  
        no shutdown  
    exit  
    no shutdown  
-----  
A:ALA-1>config>router>rsvp#
```

2.14.1 Configuring RSVP message pacing parameters

RSVP message pacing maintains a count of the messages that were dropped because the output queue for the egress interface was full.

Use the following syntax to configure RSVP parameters.

```
config>router>rsvp  
no shutdown  
msg-pacing  
period milli-seconds  
max-burst number
```

Example: RSVP message pacing configuration output

```
A:ALA-1>config>router>rsvp# info
-----
      keep-multiplier 5
      refresh-time 60
      msg-pacing
        period 400
        max-burst 400
      exit
      interface "system"
        no shutdown
      exit
      interface to-104
        hello-interval 4000
        no shutdown
      exit
      no shutdown
-----
A:ALA-1>config>router>rsvp#
```

2.14.2 Configuring graceful shutdown

Enable TE graceful shutdown on the maintenance interface using the **config>router>rsvp>interface>graceful-shutdown** command.

Disable graceful shutdown by executing the **no** form of the command at the RSVP interface level or at the RSVP level. This restores the user-configured TE parameters of the maintenance links, and the 7210 SAS maintenance node floods them.

2.15 MPLS configuration management tasks

This section describes the MPLS configuration management tasks.

2.15.1 Modifying MPLS parameters



Note:

You must shut down MPLS entities to modify parameters. Re-enable (**no shutdown**) the entity for the change to take effect.

2.15.2 Modifying an MPLS LSP

Some MPLS LSP parameters such as primary and secondary, must be shut down before they can be edited or deleted from the configuration.

Example: MPLS LSP configuration output

```
A:ALA-1>>config>router>mpls>lsp# info
-----
      shutdown
      to 10.10.10.104
-----
```

```
from 10.10.10.103
rsvp-resv-style ff
include "red"
exclude "green"
fast-reroute one-to-one
exit
primary "to-NYC"
    hop-limit 50
exit
secondary "secondary-path"
exit
-----
A:ALA-1>config>router>mpls#
```

2.15.3 Modifying MPLS path parameters

To modify path parameters, the **config>router>mpls>path** context must be shut down first.

Example: Path configuration output

```
A:ALA-1>config>router>mpls# info
#-----
echo "MPLS"
#-----
...
    path "secondary-path"
        hop 1 10.10.0.111 strict
        hop 2 10.10.0.222 strict
        hop 3 10.10.0.123 strict
        no shutdown
    exit
    path "to-NYC"
        hop 1 10.10.10.104 strict
        hop 2 10.10.0.210 strict
        no shutdown
    exit
...
-----
A:ALA-1>config>router>mpls#
```

2.15.4 Modifying MPLS static LSP parameters

To modify static LSP parameters, the **config>router>mpls>path** context must be shut down first.

Example: Static LSP configuration output

```
A:ALA-1>config>router>mpls# info
-----
...
    static-lsp "static-LSP"
        to 10.10.10.234
        push 102704 nexthop 10.10.8.114
        no shutdown
    exit
    no shutdown
-----
A:ALA-1>config>router>mpls#
```


2.15.5 Deleting an MPLS interface

To delete an interface from the MPLS configuration, the interface must be shut down first.

Use the following syntax to delete an interface from the MPLS configuration.

```
mpls
[no] interface ip-int-name
shutdown
```

```
A:ALA-1>config>router>mpls# info
-----
...
admin-group "green" 15
      admin-group "red" 25
      admin-group "yellow" 20
      interface "system"
      exit
      no shutdown
-----
A:ALA-1>config>router>mpls#
```

2.16 RSVP configuration management tasks

This section describes the RSVP configuration management tasks.

2.16.1 Modifying RSVP parameters

Only interfaces configured in the MPLS context can be modified in the RSVP context.

The **no rsvp** command deletes this RSVP protocol instance and removes all configuration parameters for this RSVP instance. The **shutdown** command suspends the execution and maintains the existing configuration.

Example

The following example displays a modified RSVP configuration example.

```
A:ALA-1>config>router>rsvp# info
-----
      keep-multiplier 5
      refresh-time 60
      msg-pacing
        period 400
        max-burst 400
      exit
      interface "system"
      exit
      interface "test1"
        hello-interval 5000
      exit
      no shutdown
-----
A:ALA-1>config>router>rsvp#
```

2.16.2 Modifying RSVP message pacing parameters

RSVP message pacing maintains a count of the messages that were dropped because the output queue for the egress interface was full.

Example

The following is a sample modified RSVP message pacing configuration output.

```
A:ALA-1>config>router>rsvp# info
-----
      keep-multiplier 5
      refresh-time 60
      msg-pacing
        period 200
        max-burst 200
      exit
      interface "system"
      exit
      interface "to-104"
      exit
      no shutdown
-----
A:ALA-1>config>router>rsvp#
```

2.16.3 Deleting an interface from RSVP

Interfaces cannot be deleted directly from the RSVP configuration. An interface must have been configured in the MPLS context and then the RSVP context. The interface must first be deleted from the MPLS context. This removes the association from RSVP.

See [Deleting an MPLS interface](#) for information about deleting an MPLS interface.

2.17 MPLS/RSVP command reference

2.17.1 Command hierarchies

- [MPLS commands](#)
- [MPLS-TP commands](#)
- [MPLS LSP commands](#)
- [MPLS-TP LSP commands](#)
- [MPLS path commands](#)
- [RSVP commands](#)
- [Show commands](#)
- [Tools commands](#)

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

2.17.1.1 MPLS commands

```

config
- router
- [no] mpls
- [no] admin-group-frr
- bypass-resignal-timer minutes
- no bypass-resignal-timer
- [no] cspf-on-loose-hop
- dynamic-bypass [enable | disable]
- [no] frr-object
- hold-timer seconds
- no hold-timer
- [no] interface ip-int-name
- [no] admin-group group-name [group-name...(up to 5 max)]
- label-map in-label
- no label-map in-label
- no pop
- pop
- no shutdown
- shutdown
- swap out-label nexthop ip-address
- swap implicit-null-label nexthop ip-address
- no swap
- no shutdown
- shutdown
- [no] srlg-group group-name [group-name...(up to 5 max)]
- te-metric metric
- no te-metric
- [no] p2mp-resignal-timer minutes
- pce-report rsvp-te {enable | disable}
- resignal-timer minutes
- no resignal-timer
- [no] shutdown
- [no] srlg-database
- [no] router-id router-addr
- [no] interface ip-addr srlg-group group-name [group-name...(up to 5 max)]
- [no] shutdown
- [no] srlg-frr [strict]
- [no] static-lsp lsp-name
- no push label
- push label nexthop ip-address
- [no] shutdown
- to ip-address
- [no] static-lsp-fast-retry seconds
- user-srlg-db [enable | disable]
- mpls-labels
- static-label-range static range
- no static-label-range

```

2.17.1.2 MPLS-TP commands



Note:

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

```

config
- router
- [no] mpls
- [no] mpls-tp
- global-id global-id
- no global-id
- node-id node-id
- no node-id
- [no] oam-template name
- bfd-template name
- no bfd-template
- hold-time-down timer
- no hold-time-down
- hold-time-up timer
- no hold-time-up
- protection-template name
- no protection-template
- rapid-psc-timer interval
- no rapid-psc-timer
- [no] revertive
- slow-psc-timer interval
- no slow-psc-timer
- wait-to-restore interval
- no wait-to-restore
- [no] shutdown
- tp-tunnel-id-range start-id end-id
- no tp-tunnel-id-range
- transit-path path-name
- no transit-path
- [no] forward-path in-label out-label out-label out-link interface
name [next-hop ip-address]
- in-label in-label out-label out-label out-link if-name [next-
hop next-hop]
- path-id {lsp-num lsp-num | working-path | protect-path [src-global-
id src-global-id] src-node-id src-node-id src-tunnel-num src-tunnel-num [dest-global-id dest-
global-id] dest-node-id dest-node-id [dest-tunnel-num dest-tunnel-num]}
- no path-id
- [no] reverse-path in-label out-label out-label out-link interface
name [next-hop ip-address]
- in-label in-label out-label out-label out-link if-name [next-
hop next-hop]
- [no] shutdown

```

2.17.1.3 MPLS LSP commands

```

config
- router
- [no] mpls
- [no] lsp lsp-name [bypass-only | p2mp-lsp | mpls-tp src-tunnel-num]
- [no] adaptive
- [no] adspec
- bgp-transport-tunnel {include | exclude}

```

```
- [no] cspf [use-te-metric]
- [no] exclude group-name [group-name...(up to 5 max)]
- fast-reroute frr-method
- no fast-reroute
  - hop-limit number
  - no hop-limit
  - [no] node-protect
  - [no] propagate-admin-group
- from ip-address
- hop-limit number
- no hop-limit
- [no] include group-name [group-name...(up to 5 max)]
- [no] ldp-over-rsvp [include | exclude]
- metric metric
- path-profile profile-id [path-group group-id]
- no path-profile profile-id
- [no] pce-computation
- [no] pce-control
- pce-report {enable | disable | inherit}
- [no] propagate-admin-group
- [no] to id
- [no] primary path-name
  - [no] adaptive
  - bandwidth rate-in-mps
  - no bandwidth
  - [no] exclude group-name [group-name...(up to 5 max)]
  - hop-limit number
  - no hop-limit
  - [no] include group-name [group-name...(up to 5 max)]
  - [no] record
  - [no] record-label
  - [no] shutdown
- retry-limit number
- no retry-limit
- retry-timer seconds
- no retry-timer
- rsvp-resv-style {se | ff}
- [no] secondary path-name
  - [no] adaptive
  - bandwidth rate-in-mps
  - no bandwidth
  - [no] exclude group-name [group-name...(up to 5 max)]
  - hop-limit number
  - no hop-limit
  - [no] include group-name [group-name...(up to 5 max)]
  - [no] path-preference
  - [no] record
  - [no] record-label
  - [no] shutdown
  - [no] srlg
  - [no] standby
- [no] shutdown
- to ip-address
- vprn-auto-bind [include | exclude]
- lsp-template template-name p2mp
- no lsp-template template-name
  - cspf [use-te-metric]
  - no cspf
  - default-path path-name
  - exclude group-name [group-name...(up to 5 max)]
  - no exclude [group-name [group-name...(up to 5 max)]]
  - include group-name [group-name...(up to 5 max)]
  - no include [group-name [group-name...(up to 5 max)]]
  - [no] propagate-admin-group
```

```
- [no] record
- [no] record-label
- retry-limit number
- no retry-limit
- retry-timer seconds
- no retry-timer
- [no] shutdown
```

2.17.1.4 MPLS-TP LSP commands



Note:

MPLS-TP LSP commands are only supported on 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T (network mode).

```
config
- router
- [no] mpls
- [no] lsp lsp-name [bypass-only | mpls-tp src-tunnel-num]
- [no] protect-tp-path
- in-label in-label
- lsp-num lsp-number
- no lsp-num
- [no] mep
- bfd-enable [bfd-mode]
- no bfd-enable
- oam-template [32 chars max]
- no oam-template
- protection-template [256 chars max]
- no protection-template
- [no] shutdown
- out-label out-label out-link if-name [next-hop ip-address]
- no out-label
- [no] shutdown
- [no] working-tp-path
- in-label in-label
- lsp-num lsp-number
- no lsp-num
- [no] mep
- bfd-enable [bfd-mode]
- no bfd-enable
- oam-template name
- no oam-template
- [no] shutdown
- out-label out-label out-link if-name [next-hop ip-address]
- no out-label
- [no] shutdown
```

2.17.1.5 MPLS path commands

```
config
- router
- [no] mpls
- [no] path path-name
- hop hop-index ip-address {strict | loose}
- no hop hop-index
- [no] shutdown
- [no] static-lsp lsp-name
```

```
- push label nexthop ip-address
- no push out-label
- to ip-addr
- [no] shutdown
```

2.17.1.6 RSVP commands

```
config
- router
  - [no] rsvp
    - [no] graceful-shutdown
    - [no] implicit-null-label
    - [no] interface ip-int-name
      - authentication-key [authentication-key | hash-key] [hash | hash2]
      - no authentication-key
      - [no] bfd-enable
      - [no] graceful-shutdown
      - hello-interval milli-seconds
      - no hello-interval
      - [no] refresh-reduction
        - [no] reliable-delivery
      - [no] shutdown
      - subscription percentage
      - no subscription
    - keep-multiplier number
    - no keep-multiplier
    - node-id-in-rro {include | exclude}
    - [no] msg-pacing
      - max-burst number
      - no max-burst
      - period milli-seconds
      - no period
    - rapid-retransmit-time hundred-milliseconds
    - no rapid-retransmit-time
    - rapid-retry-limit number
    - no rapid-retry-limit
    - refresh-time seconds
    - no refresh-time
    - [no] shutdown
```

2.17.1.7 Show commands

```
show
- router
  - mpls
    - admin-group group-name
    - bypass-tunnel [to ip-address] [protected-lsp name] [dynamic | manual | p2mp]
  [detail]
    - interface [ip-int-name | ip-address] [label-map label]
    - interface [ip-int-name | ip-address]
    - lsp [lsp-name] [status {up | down}] [from ip-address | to ip-address] [detail]
    - lsp {transit | terminate} [status {up | down}] [from ip-address | to ip-
address | lsp-name name] [detail]
    - lsp count
    - lsp lsp-name activepath
    - lsp [lsp-name] path [path-name] [status {up | down}] [detail]
    - mpls-labels
      - label start-label [end-label | in-use | label-owner]
```

```

- label-range
- mpls-tp
- oam-template [template-name] [associations]
- protection-template [template-name] [associations]
- status
- transit-path [path-name] [detail]
- path [path-name] [lsp-binding]
- srlg-database [router-id ip-address] [interface ip-address]
- srlg-group [group-name]
- static-lsp [lsp-name]
- static-lsp {transit | terminate}
- static-lsp count
- status

show
- router
- rsvp
- interface [interface [ip-int-name]] statistics [detail]
- neighbor [ip-address] [detail]
- session [session-type] [from ip-address | to ip-address | lsp-name name] [status
{up | down}] [detail]
- statistics
- status

```

2.17.1.8 Tools commands

```

- perform
- router
- mpls
- cspf to ip-addr [from ip-addr] [bandwidth bandwidth] [include-bitmap bitmap]
[exclude-bitmap bitmap] [hop-limit limit] [exclude-address excl-addr [excl-addr...(up to 8
max)]] [use-te-metric] [strict-srlg] [srlggroup
grp-id...(up to 8 max)] [skip-interface interface-name]
- resignal {lsp lsp-name path path-name | delay minutes}
- resignal-bypass {lsp bypass-lsp-name [force] | delay minutes}
- tp-tunnel
- clear id tunnel-id
- clear lsp-name
- force id tunnel-id
- force lsp-name
- lockout id tunnel-id
- lockout lsp-name
- manual id tunnel-id
- manual lsp-name
- trap-suppress number-of-traps time-interval

```

2.17.1.9 Clear commands

```

clear
- router
- mpls
- interface [ip-int-name]
- lsp lsp-name
- rsvp
- interface [ip-int-name] [statistics]
- statistics

```


2.17.1.10 Debug commands

```
debug
- router
- mpls [lsp lsp-name] [sender source-address] [endpoint endpoint-address] [tunnel-
id tunnel-id] [lsp-id lsp-id]
- no mpls
- [no] event
- all [detail]
- no all
- frr [detail]
- no frr
- iom [detail]
- no iom
- lsp-setup [detail]
- no lsp-setup
- mbb [detail]
- no mbb
- misc [detail]
- no misc
- xc [detail]
- no xc
- rsvp [lsp lsp-name] [sender source-address] [endpoint endpoint-address] [tunnel-
id tunnel-id] [lsp-id lsp-id] [interface ip-int-name]
- no rsvp
- [no] event
- all [detail]
- no all
- auth
- no auth
- misc [detail]
- no misc
- nbr [detail]
- no nbr
- path [detail]
- no path
- resv [detail]
- no resv
- rr
- no rr
- [no] packet
- all [detail]
- no all
- ack
- bundle [detail]
- no bundle
- hello [detail]
- no hello
- path [detail]
- no path
- patherr [detail]
- no patherr
- pathtear [detail]
- no pathtear
- resv [detail]
- no resv
- resvrr [detail]
- no resvrr
- resvtear [detail]
- no resvtear
- srefresh [detail]
```

- no [srefresh](#)

2.17.2 Command descriptions

- [MPLS configuration commands](#)
- [RSVP configuration commands](#)
- [Show commands](#)
- [Tools commands](#)
- [Clear commands](#)
- [Debug commands](#)

2.17.2.1 MPLS configuration commands

- [Generic commands](#)
- [MPLS commands](#)
- [MPLS label commands](#)
- [MPLS interface commands](#)
- [MPLS-TP commands](#)
- [LSP commands](#)
- [Primary and secondary path commands](#)
- [LSP path commands](#)
- [Static LSP commands](#)

2.17.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

config>router>mpls

config>router>mpls>interface

config>router>mpls>lsp>primary

config>router>mpls>lsp>secondary

config>router>mpls>lsp-template

Platforms

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

Description



Note:

The **config>router>mpls>lsp-template** context is only supported on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T (operating in network mode), and 7210 SAS-Sx/S 1/10GE (operating in standalone and standalone-VC mode).

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

MPLS is not enabled by default and must be explicitly enabled (**no shutdown**).

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

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

Default

no shutdown

2.17.2.1.2 MPLS commands

mpls

Syntax

[no] mpls

Context

config>router

Platforms

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

Description

Commands in this context configure MPLS parameters. MPLS is not enabled by default and must be explicitly enabled (**no shutdown**). The **shutdown** command administratively disables MPLS.

The **no** form of this command deletes this MPLS protocol instance; this will remove all configuration parameters for this MPLS instance.

MPLS must be **shutdown** before the MPLS instance can be deleted. If MPLS is not **shutdown**, when the **no mpls** command is executed, a warning message is displayed on the console indicating that MPLS is still administratively up.

admin-group-frr

Syntax

[no] admin-group-frr

Context

config>router>mpls

Platforms

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

Description

This command enables the use of admin-group constraints when a manual or dynamic bypass LSP is associated with the primary LSP path at a Point-of-Local Repair (PLR) node.

When this command is enabled, each PLR node reads the admin-group constraints in the FAST_REROUTE object included in the Path message of the LSP primary path. If the object is not included, the PLR reads the Session Attribute object in the Path message.

If the PLR is also the ingress LER for the LSP primary path, only the admin-group constraints from the LSP or path level configurations are used.

Next, the PLR node uses the admin-group and other constraints, such as hop-limit and SRLG, to select a manual or dynamic bypass LSP among the bypass LSPs that are already in use.

If none of the manual or dynamic bypass LSPs satisfy the admin-group and other constraints, the PLR node requests the CSPF for a path that merges the closest to the protected link or node and that includes or excludes the specified admin-group IDs.

Modifying the configuration of this command does not affect existing bypass associations. The change only applies to new attempts to find a valid bypass.

The **no** form of this command disables the use of administrative group constraints on an FRR backup LSP at a PLR node.

Default

no admin-group-frr

bypass-resignal-timer

Syntax

bypass-resignal-timer *minutes*

no bypass-resignal-timer

Context

config>router>mpls

Platforms

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

Description

This command triggers the periodic global reoptimization of all dynamic bypass LSP paths associated with RSVP P2P LSP. The operation is performed at each expiry of the user-configurable bypass LSP resignal timer.

When this command is enabled, MPLS requests the CSPF for the best path for each dynamic bypass LSP originated on this node. The constraints of the first associated LSP primary path that originally triggered the signaling of the bypass LSP must be satisfied. To do this, MPLS saves the original Path State Block (PSB) of that LSP primary path, even if the latter is torn down.

If CSPF returns no path or returns a new path with a cost that is higher than the current path, MPLS does not signal the new bypass path. If CSPF returns a new path with a cost that is lower than the current path, MPLS signals it. Also, if the new bypass path is SRLG strict disjoint with the primary path of the original PSB while the current path is SLRG loose disjoint, the manual bypass path is resigned, regardless of cost comparison.

After the new path is successfully signaled, MPLS evaluates each PSB of each PLR (that is, each unique avoid-node or avoid-link constraint) associated with the older bypass LSP path to check if the corresponding LSP primary path constraints are still satisfied by the new bypass LSP path. If so, the PSB association is moved to the new bypass LSP.

Each PSB for which constraints are not satisfied remains associated with the older bypass LSP and is checked at the next background PSB re-evaluation, or at the next timer or manual bypass reoptimization. If the older bypass LSP is SRLG disjoint with a primary path that has the non-strict SRLG constraint while the new bypass LSP is not SRLG disjoint, the PSB association is not moved.

If a specific PLR associated with a bypass LSP is active, the corresponding PSBs remain associated with the older bypass LSP until the global revertive Make-Before-Break (MBB) tears down all corresponding primary paths, which will also cause the older bypass LSP to be torn down.

This periodic bypass reoptimization feature also implements a background PSB re-evaluation task that audits in the background each RSVP session and determines if an existing manual or dynamic bypass is more optimal for that session. If so, it moves the PSB association to this existing bypass. If the PLR for this session is active, no action is taken and the PSB is re-examined at the next re-evaluation.

The periodic bypass reoptimization feature evaluates only the PSBs of the PLRs associated with that bypass LSP and only against the new bypass LSP path. The background re-evaluation task, however, audits all PSBs on the system against all existing manual and dynamic bypass LSPs.

PSBs that have not been moved by the dynamic or manual reoptimization of a bypass LSP because the PSB constraints have not been met by the new signaled bypass LSP path are re-evaluated by the background task against all existing manual and dynamic bypass LSPs.

Finally, the background re-evaluation task checks for PSBs that have requested node-protect bypass LSP but are currently associated with a link-protect bypass LSP, and PSBs that requested FRR protection but have no association. This is in addition to the attempt made at the receipt of a Resv message on the protected LSP path to accelerate the association.

This feature is not supported with inter-area dynamic bypass LSP and bypass LSP protecting S2L paths of a P2MP LSP.

The **no** form of this command disables the periodic global reoptimization of dynamic bypass LSP paths.

Default

no bypass-resignal timer

Parameters

minutes

Specifies the time, in minutes, that MPLS waits before attempting to resignal dynamic bypass LSP paths originated on the system.

Values 30 to 10080

cspf-on-loose-hop

Syntax

[no] cspf-on-loose-hop

Context

config>router>mpls

Platforms

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

Description

This command enables the option to perform CSPF calculations until the next loose hop or the final destination of the LSP on LSR. On receiving a PATH message on the LSR and processing all local hops in the received ERO, if the next hop is loose, the LSR node does a CSPF calculation until the next loose hop. On successful completion of the CSPF calculation, the ERO in the PATH message is modified to include newly calculated intermediate hops and the message is propagated forward to the next hop. This allows for the setting up of inter-area LSPs based on the ERO expansion method.



Note:

The LSP may fail to set up if this command is enabled on an LSR that is not an area border router and that receives a PATH message without a proper next loose hop in the ERO. The **cspf-on-loose-hop** command can change dynamically and is applied to the new LSP setup after changes are made.

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

Default

no cspf-on-loose-hop

dynamic-bypass

Syntax

dynamic-bypass [enable | disable]

no dynamic-bypass

Context

config>router>mpls

Platforms

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

Description

This command disables the creation of dynamic bypass LSPs in FRR. One or more manual bypass LSPs must be configured to protect the primary LSP path at the PLR nodes.

If the 7210 SAS is used as an egress LER and is a merge point, implicit null must be enabled for use of manual bypass or dynamic bypass (FRR facility).

Default

dynamic-bypass enable

frr-object

Syntax

[no] frr-object

Context

config>router>mpls

Platforms

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

Description

This command configures whether fast reroute for LSPs using the facility bypass method is signaled with or without the fast reroute object using the **one-to-one** keyword. The value is ignored if fast reroute is disabled for the LSP or if the LSP is using one-to-one backup.

Default

frr-object

hold-timer

Syntax

hold-timer *seconds*

no hold-timer

Context

config>router>mpls

Platforms

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

Description

This command specifies the amount of time that the ingress node waits before programming its data plane and declaring to the service module that the LSP state is up.

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

Default

hold-timer 1

Parameters

seconds

Specifies the hold time, in seconds.

Values 0 to 10

p2mp-resignal-timer

Syntax

p2mp-resignal-timer *minutes*

no p2mp-resignal-timer

Context

config>router>mpls

Platforms

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

Description

This command configures the resignal timer for a P2MP LSP instance.

MPLS requests CSPF to recompute the whole set of S2L paths of a specific active P2MP instance each time the P2MP resignal timer expires. The P2MP resignal timer is configured separately from the P2P LSP parameter. MPLS performs a global MBB and moves each S2L sub-LSP in the instance into its new path using a new P2MP LSP ID if the global MBB is successful, regardless of the cost of the new S2L path.

The **no** form of this command disables the timer-based resignaling of P2MP LSPs on this system.

Default

no resignal-timer

Parameters

minutes

Specifies the time, in minutes, that MPLS waits before attempting to resignal the P2MP LSP instance.

Values 60 to 10080

pce-report

Syntax

pce-report rsvp-te {enable | disable}

Context

config>router>mpls

Platforms

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

Description

This command configures the reporting mode for RSVP-TE LSPs.

The PCC LSP database is synchronized with the PCE LSP database using the PCEP PCRpt (PCE Report) message for PCC-controlled, PCE-computed, and PCE-controlled LSPs.

This global MPLS-level **pce-report** command enables or disables PCE reporting for all RSVP-TE LSPs during PCE LSP database synchronization. The PCC reports both CSPF and non-CSPF LSPs.

The LSP-level **pce-report** command (**config>router>mpls>lsp>pce-report**) overrides the global configuration for reporting an LSP to the PCE. The default configuration, which inherits the global MPLS-level configuration, is disabled (**pce-report rsvp-te disable**).

The default configuration controls the introduction of a PCE into an existing network and allows the operator to decide whether all RSVP-TE LSPs should be reported. If PCE reporting for an LSP is disabled, either because of inheritance of the global MPLS configuration or because of LSP-level configuration, enabling the **pce-control** option for the LSP has no effect.

Default

pce-report rsvp-te disable

Parameters

rsvp-te {enable | disable}

Enables or disables PCE reporting for all RSVP-TE LSPs.

resignal-timer

Syntax

resignal-timer minutes

no **resignal-timer**

Context

config>router>mpls

Platforms

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

Description

This command specifies the value for the LSP resignal timer. The resignal timer is the wait time, in minutes, before the software attempts to resignal the LSPs.

When the resignal timer expires, if the new computed path for an LSP has a better metric than the current recorded hop list, an attempt is made to resignal that LSP using the make-before-break mechanism. If the attempt to resignal an LSP fails, the LSP continues to use the existing path and a resignal is attempted the next time the timer expires.

The **no** form of this command disables timer-based LSP resignaling.

Default

no resignal-timer

Parameters

minutes

Specifies the time, in minutes, that the software waits before attempting to resignal the LSPs.

Values 30 to 10080

srlg-frr

Syntax

srlg-frr [strict]

no **srlg-frr**

Context

config>router>mpls

Platforms

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

Description

This command enables the use of the Shared Risk Link Group (SRLG) constraint in the computation of an FRR bypass or detour LSP for any primary LSP path on this system.

When this option is enabled, CSPF includes the SRLG constraint in the computation of an FRR detour or bypass for protecting the primary LSP path.

CSPF prunes all links with interfaces that belong to the same SRLG as the interface that is being protected, where the interface being protected is the outgoing interface at the PLR used by the primary path.

If one or more paths are found, the MPLS/RSVP task selects one path based on best cost and signals the setup of the FRR bypass or detour LSP. If no path is found and the user included the **strict** option, the FRR bypass or detour LSP is not setup and the MPLS/RSVP task will keep retrying the request to CSPF. If no path is found and the strict option is disabled, if a path exists that meets all the TE constraints except the SRLG constraint, the FRR bypass or detour LSP is set up.

An FRR bypass or detour LSP path is not guaranteed to be SRLG disjoint from the primary path. This is because only the SRLG constraint of the outgoing interface at the PLR that the primary path is using is checked.

When the MPLS/RSVP task is searching for a SRLG bypass tunnel to associate with the primary path of the protected LSP, the task first checks if any configured manual bypass LSP with CSPF enabled satisfies the SLRG constraints. The MPLS/RSVP skips any non-CSPF bypass LSP in the search as there is no ERO returned to check the SLRG constraint. If no path is found, the task checks if an existing dynamic bypass LSP satisfies the SLRG and other primary path constraints. If not, it will make a request to CSPF.

After the primary path of the LSP is set up and is operationally up, any subsequent changes to the SRLG group membership of an interface that the primary path is using is not be considered by the MPLS/RSVP task at the PLR for FRR bypass or detour LSP association until the next opportunity that the primary path is resigaled. The path may be resigaled because of a failure or to a make-before-break operation. Make-before-break occurs as a result of a global revertive operation, a timer based or manual reoptimization of the LSP path, or a user change to any of the path constraints.

After the FRR bypass or detour LSP path is setup and is operationally up, any subsequent changes to the SRLG group membership of an interface that the FRR bypass or detour LSP path is using would not be considered by the MPLS/RSVP task at the PLR until the next opportunity the association with the primary LSP path is rechecked. The association is rechecked if the bypass path is reoptimized. Detour paths are not reoptimized and are resigaled if the primary path is down.

Enabling or disabling **srlg-frr** only takes effect after LSP paths are resigaled. This can be achieved by shutting down and re-enabling MPLS. Another option is using the **tools>perform>router>mpls>resignal** command. Though the latter has less impact on service, only originating LSPs can be resigaled with the **tools** command. If local transit and bypass LSPs are also to be resigaled, the **tools** command must be executed on all ingress nodes in the network. The same can be locally achieved by disabling and enabling using the **configure>router>mpls>dynamic-bypass** command, but this can trigger the LSP to go down and traffic loss to occur in case detour or bypass LSP is in use.

An RSVP interface can belong to a maximum of 64 SRLG groups. The user configures the SRLG groups using the **config>router>mpls>srlg-group** command. The user associates the SRLG with an RSVP interface using the **srlg-group** command in the **config>router>mpls>interface** context.

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

Default

no srlg-frr

Parameters

strict

Specifies the name of the SRLG group within a virtual router instance.

Values	no slr-frr (default)
	srlg-frr (non-strict)
	srlg-frr strict (strict)

user-srlg-db

Syntax

user-srlg-db [enable | disable]

Context

config>router>mpls

Platforms

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

Description

This command enables the use of CSPF by the user SRLG database. When the MPLS module makes a request to CSPF for the computation of an SRLG secondary path, CSPF queries the local SRLG and computes a path after pruning links that are members of the SRLG IDs of the associated primary path. When MPLS makes a request to CSPF for an FRR bypass or detour path to associate with the primary path, CSPF queries the user SRLG database and computes a path after pruning links that are members of the SRLG IDs of the PLR outgoing interface.

If an interface is not entered into the user SRLG database, it is assumed that it does not have any SRLG membership. CSPF will not query the TE database for IGP advertised interface SRLG information.

The **disable** keyword disables the use of the user SRLG database. CSPF resumes queries into the TE database for SRLG membership information. The user SRLG database is maintained.

Default

user-srlg-db disable

Parameters

enable

Keyword to enable the use of the user SRLG database.

disable

Keyword to disable the use of the user SRLG database.

srlg-database

Syntax

[no] srlg-database

Context

config>router>mpls

Platforms

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

Description

Commands in this context manually enter the link members of SRLG groups for the entire network at any node that needs to signal LSP paths (for example, a head-end node).

The **no** form of this command deletes the entire SRLG database. CSPF assumes all interfaces have no SRLG membership association if the database was not disabled with the command **config>router>mpls>user-srlg-db disable**.

router-id

Syntax

[no] router-id *ip-address*

Context

config>router>mpls>srlg-database

Platforms

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

Description

This command manually enters the link members of SRLG groups for a specific router in the network. The user must also use this command to enter the local interface SRLG membership into the user SRLG database. Use by CSPF of all interface SRLG membership information of a specific router ID may be temporarily disabled by shutting down the node. If this occurs, CSPF will assume these interfaces have no SRLG membership association.

The **no** form of this command deletes all interface entries under the router ID.

Parameters

ip-address

Specifies the router ID for this system. This value must be the router ID configured under the base router instance, the base OSPF instance, or the base IS-IS instance.

Values [a.b.c.d]

interface

Syntax

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

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

Context

config>router>mpls>srlg-database>router-id

Platforms

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

Description

This command configures SRLG membership information for any link in the network, including links on this node, in the user SRLG database.

An interface can be associated with up to five SRLG groups for each execution of this command. The operator can associate an interface with up to 64 SRLG groups by executing the command multiple times.

CSPF will not use entered SRLG membership if an interface is not validated as part of a router ID in the routing table.

The **no** form of this command deletes a specific interface entry in this user SRLG database. The group name must already exist in the **config>router>mpls>srlg-group** context.

Parameters

ip-int-name

Specifies the name of the network IP interface. An interface name cannot be in the form of an IP address.

srlg-group group-name

Specifies the SRLG group name, up to 32 characters. Up to 1024 group names can be defined in the **config>router>mpls** context. The SRLG group names must be identical across all routers in a single domain.

label-map

Syntax

[no] **label-map** *in-label*

Context

config>router>mpls>interface

Platforms

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

Description

This command is used on transit routers when a static LSP is defined. The static LSP on the ingress router is initiated using the **config>router>mpls>static-lsp** */sp-name* command. The *in-label* is associated with either a **pop** or a **swap** action, but not both. If both actions are specified, the last action specified takes effect.

The **no** form of this command deletes the static LSP configuration associated with the *in-label*.

Parameters

in-label

Specifies the incoming MPLS label on which to match.

Values 32 to 1023

pop

Syntax

[no] pop

Context

config>router>mpls>if>label-map

Platforms

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

Description

This command specifies that the incoming label must be popped (removed). No label stacking is supported for a static LSP. The service header follows the top label. After the label is popped, the packet is forwarded based on the service header.

The **no** form of this command removes the **pop** action for the *in-label*.

shutdown

Syntax

[no] shutdown

Context

config>router>mpls>if>label-map

Platforms

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

Description

This command disables the label map definition. This drops all packets that match the *in-label* specified in the **label-map** *in-label* command.

The **no** form of this command administratively enables the defined label map action.

Default

no shutdown

swap

Syntax

swap {*out-label* | *implicit-null-label*} *nexthop ip-address*

no swap {*out-label* | *implicit-null-label*}

Context

config>router>mpls>interface>label-map

Platforms

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

Description

This command swaps the incoming label and specifies the outgoing label and next hop IP address on an LSR for a static LSP.

The **no** form of this command removes the swap action associated with the *in-label*.

Parameters

implicit-null-label

Keyword to specify the use of the implicit label value for the outgoing label of the swap operation.

out-label

Specifies the label value to be swapped with the *in-label*. Label values 16 through 1,048,575 are defined as follows.

- Label values 16 through 31 are reserved.
- Label values 32 through 1,023 are available for static assignment.
- Label values 1,024 through 2,047 are reserved for future use.
- Label values 2,048 through 18,431 are statically assigned for services.
- Label values 28,672 through 131,071 are dynamically assigned for both MPLS and services.

- Label values 131,072 through 1,048,575 are reserved for future use.

Values 16 to 1048575

nexthop *ip-address*

Specifies the IP address to forward to. If an ARP entry for the next hop exists, the static LSP is marked operational. If an ARP entry does not exist, the software sets the operational status of the static LSP to down and continues to the ARP for the configured next hop. Software will continue to the ARP for the configured next hop at a fixed interval.

static-lsp

Syntax

[no] **static-lsp** *lsp-name*

Context

config>router>mpls

Platforms

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

Description

This command configures a static LSP on the ingress router. The static LSP is a manually setup LSP where the next-hop IP address and the outgoing label (push) must be specified.

The LSP must first be shut down to delete it. If the LSP is not shut down, the **no static-lsp *lsp-name*** command generates a warning message on the console indicating that the LSP is administratively up.

The **no** form of this command deletes this static LSP and associated information.

Parameters

lsp-name

Specifies the LSP name, up to 32 characters.

push

Syntax

no push label

push label nexthop *ip-address*

Context

config>router>mpls>static-lsp

Platforms

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

Description

This command specifies the label to be pushed on the label stack and the next hop IP address for the static LSP.

The **no** form of this command removes the association of the label to push for the static LSP.

Parameters

label

Specifies the label to push on the label stack. Label values 16 through 1,048,575 are defined as follows.

- Label values 16 through 31 are reserved.
- Label values 32 through 1,023 are available for static assignment.
- Label values 1,024 through 2,047 are reserved for future use.
- Label values 2,048 through 18,431 are statically assigned for services.
- Label values 28,672 through 131,071 are dynamically assigned for both MPLS and services.
- Label values 131,072 through 1,048,575 are reserved for future use.

Values 16 to 1048575

nexthop ip-address

Specifies the IP address of the next hop toward the LSP egress router. If an ARP entry for the next hop exists, the static LSP is marked operational.

If an ARP entry does not exist, software sets the operational status of the static LSP to down and continues to ARP for the configured next hop. Software continuously tries to ARP for the configured next hop at a fixed interval.

shutdown

Syntax

[no] shutdown

Context

config>router>mpls>static-lsp

Platforms

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

Description

This command administratively disables the static LSP.

The **no** form of this command administratively enables the static LSP.

Default

shutdown

to

Syntax

to *ip-address*

Context

config>router>mpls>static-lsp

Platforms

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

Description

This command configures the system IP address of the egress router for the static LSP. This command is required while creating an LSP. For LSPs that are used as transport tunnels for services, the to IP address must be the system IP address. If the to address does not match the SDP address, the LSP is not included in the SDP definition.

Parameters

ip-address

Specifies the system IP address of the egress router.

Values a.b.c.d

static-lsp-fast-retry

Syntax

static-lsp-fast-retry *seconds*

[no] static-lsp-fast-retry

Context

config>router>mpls

Platforms

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

Description

This command configures the fast retry timer value for a static LSP.

When a static LSP is trying to come up, the MPLS request for the ARP entry of the LSP next hop may fail when it is made while the next hop is still down or unavailable. In that case, MPLS starts a retry timer before making the next request. This enhancement allows the user to configure the retry timer so that the LSP comes up as soon as the next hop is up.

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

Default

no static-fast-retry-timer

Parameters

seconds

Specifies the value, in seconds, used as the fast retry timer for a static LSP.

Values 1 to 30

2.17.2.1.3 MPLS label commands

mpls-labels

Syntax

mpls-labels

Context

config>router

Platforms

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

Description

Commands in this context configure MPLS labels on the ingress router.

static-label-range

Syntax

static-label-range *static-range*

no static-label-range

Context

config>router>mpls-labels

Platforms

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

Description

This command configures the static label range on the ingress router.

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

Default

static-label-range 18400

Parameters

static-range

Specifies the static label range.

Values 0 to 131040

2.17.2.1.4 MPLS interface commands

interface

Syntax

[no] interface *ip-int-name*

Context

config>router>mpls

Platforms

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

Description

This command specifies MPLS protocol support on an IP interface. MPLS commands are not executed on an IP interface where MPLS is not enabled. An MPLS interface must be explicitly enabled (**no shutdown**).

The **no** form of this command deletes all MPLS commands, such as **label-map**, that are defined under the interface. The MPLS interface must first be shut down to delete the interface definition. If the interface is not shut down, the **no interface ip-int-name** command does nothing except issue a warning message on the console indicating that the interface is administratively up.

Default

shutdown

Parameters

ip-int-name

Specifies the name of the network IP interface, up to 32 characters. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

admin-group

Syntax

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

Context

config>router>mpls>interface

Platforms

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

Description

This command configures administrative groups that this interface supports.

This information is advertised as part of OSPF and IS-IS to help CSPF compute constrained LSPs that must include or exclude specific administrative groups. An MPLS interface is assumed to belong to all the administrative groups unless the **admin-group** command is issued under the interface configuration. When the **admin-group** command is issued, the interface is assumed to belong to only the specifically listed groups for that command.

Each single operation of the **admin-group** command allows a maximum of five groups to be specified at a time. However, a maximum of 32 groups can be specified per interface through multiple operations.

Default

no admin-group

Parameters

group-name

Specifies the name of the group, up to 32 characters. The group names should be the same across all routers in the MPLS domain.

srfg-group

Syntax

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

Context

config>router>mpls>interface

Platforms

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

Description

This command associates an RSVP interface to SRLG groups. An interface can belong to up to 64 SRLG groups. However, each single operation of the **srlg-group** command allows a maximum of five groups to be specified at a time.

The **no** form of this command deletes the association of the interface to the SRLG group.

Parameters

group-name

Specifies the name of the SRLG group, up to 32 characters, within a virtual router instance.

te-metric

Syntax

te-metric value

no te-metric

Context

config>router>mpls>interface

Platforms

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

Description

This command configures the traffic engineering metric used on the interface. This metric is in addition to the interface metric used by IGP for the shortest path computation.

This metric is flooded as part of the TE parameters for the interface using an opaque LSA or an LSP. The IS-IS TE metric is encoded as sub-TLV 18 as part of the extended IS reachability TLV, and the metric value is encoded as a 24-bit unsigned integer. The OSPF TE metric is encoded as a sub-TLV Type 5 in the Link TLV, and the metric value is encoded as a 32-bit unsigned integer.

When the use of the TE metric is enabled for an LSP, CSPF first prunes all links in the network topology that do not meet the constraints specified for the LSP path. Such constraints include bandwidth, admin-groups, and hop limit. Then, CSPF runs an SPF on the remaining links. The shortest path among all the SPF paths will be selected based on the TE metric instead of the IGP metric, which is used by default.

The TE metric in CSPF LSP path computation can be configured using the **config>router>mpls>lsp>cspf>use-te-metric** CLI command.

The TE metric is only used in CSPF computations for MPLS paths and not in the regular SPF computation for IP reachability. The value of the IGP metric is advertised in the TE metric sub-TLV by IS-IS and OSPF.

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

Default

no te-metric

Parameters

value

Specifies the metric value.

Values 1 to 16777215

2.17.2.1.5 MPLS-TP commands

`mpls-tp`

Syntax

`[no] mpls-tp`

Context

`config>router>mpls`

Platforms

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

Description

Commands in this context configure generic MPLS-TP parameters and MPLS-TP transit paths. If a user configures **no mpls**, normally the entire MPLS configuration is deleted. However, in the case of **mpls-tp**, a check is made that there is no other **mpls-tp** configuration (for example, services or LSPs using MPLS TP on the node). The **mpls-tp** context cannot be deleted if MPLS-TP LSPs or SDPs exist on the system.

A shutdown of **mpls-tp** will bring down all MPLS-TP LSPs on the system.

Default

`no mpls-tp`

`tp-tunnel-id-range`

Syntax

`tp-tunnel-id-range start-id end-id`

`no tp-tunnel-id-range`

Context

`config>router>mpls>mpls-tp`

Platforms

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

Description

This command configures the range of MPLS tunnel IDs reserved for MPLS-TP LSPs. The maximum difference between the *start-id* and *end-id* is 4000.

The tunnel ID is the RSVP-TE tunnel ID. This maps to the MPLS-TP tunnel number. In some cases, dynamic LSPs may cause fragmentation to the number space such that the contiguous range [*end-id* – *start-id*] is not available. In these cases, the command fails.

There are no default values for the *start-id* and *end-id* of the tunnel ID range, and they must be configured to enable MPLS-TP.

Default

no tunnel-id-range

Parameters

start-id

Specifies the start ID.

Values 1 to 61440

end-id

Specifies the end ID.

Values 1 to 61440

oam-template

Syntax

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

Context

config>router>mpls>mpls-tp

Platforms

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

Description

This command enables or edits an OAM template context. Generally, applicable proactive OAM parameters are configured using templates. The top-level template is the OAM template.

Generic MPLS-TP OAM and fault management parameters are configured in the OAM template.

Proactive CC/CV uses BFD and parameters such as Tx/Rx timer intervals, multiplier, and other session or fault management parameters specific to BFD that are configured using a BFD template, which is referenced from the OAM template.

Default

no oam-template

Parameters

name

Specifies a text string name for the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes. Named OAM templates are referenced from the MPLS-TP path MEP configuration.

hold-time-down

Syntax

hold-time-down *timer*

no hold-time-down

Context

config>router>mpls>mpls-tp>oam-template

Platforms

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

Description

This command configures the hold-down dampening timer. It is equivalent to a hold-off timer.

Default

no hold-time-down

Parameters

interval

Specifies the hold-down dampening timer interval.

Values 0 to 5000 deciseconds in 10 ms increments

hold-time-up

Syntax

hold-time-up *timer*

no hold-time-up

Context

config>router>mpls>mpls-tp>oam-template

Platforms

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

Description

This command configures the hold-up dampening timer. This can be used to provide additional dampening to the state of proactive CC BFD sessions.

Default

no hold-time-up

Parameters

interval

Specifies the hold-up dampening timer interval.

Values 0 to 500 deciseconds, in 100 ms increments

Default 2 seconds

bfd-template

Syntax

bfd-template name

no bfd-template

Context

config>router>mpls>mpls-tp>oam-template

Platforms

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

Description

This command configures a named BFD template to be referenced by an OAM template.

Default

no bfd-template

Parameters

name

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

protection-template

Syntax

protection-template name

no protection-template

Context

config>router>mpls>mpls-tp

Platforms

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

Description

This command creates or edits a named protection template context. Protection templates are used to define generally applicable protection parameters for MPLS-TP tunnels. Only linear protection is supported; the application of a named template to an MPLS-TP LSP implies that linear protection is used. A protection template is applied under the MEP context of the protect-path of an MPLS-TP LSP.

Default

no protection-template

Parameters

name

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

revertive

Syntax

[no] revertive

Context

config>router>mpls>mpls-tp>protection-template

Platforms

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

Description

This command configures revertive behavior for MPLS-TP linear protection. The protect-tp-path MEP must be in the **shutdown** state for the MPLS-TP LSPs referencing this protection template to change the revertive parameter.

Default

revertive

wait-to-restore

Syntax

wait-to-restore interval

no wait-to-restore

Context

config>router>mpls>mpls-tp>protection-template

Platforms

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

Description

This command configures the WTR timer. It determines how long to wait until the active path of an MPLS-TP LSP is restored to the working path following the clearing of a defect on the working path. It is applicable only for revertive mode.

Default

no wait-to-restore

Parameters

interval

Specifies the WTR timer interval.

Values 0 to 720 seconds, in 1 second increments

rapid-psc-timer

Syntax

rapid-psc-timer interval

no rapid-psc-timer

Context

config>router>mpls>mpls-tp>protection-template

Platforms

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

Description

This command configures the rapid timer value used for protection switching coordination (PSC) packets for MPLS-TP linear protection, in accordance with RFC 6378.

Default

no rapid-psc-timer

Parameters

interval

Specifies the rapid timer interval, in milliseconds.

Values 10, 100, 1000

Default 10

slow-psc-timer

Syntax

slow-psc-timer interval

no slow-psc-timer

Context

config>router>mpls>mpls-tp>protection-template

Platforms

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

Description

This command configures the slow timer value used for PSC packets for MPLS-TP linear protection, in accordance with RFC 6378.

Default

no rapid-psc-timer

Parameters

interval

Specifies the slow timer interval, in milliseconds.

Values 10, 100, 1000

global-id

Syntax

global-id global-id

no global-id

Context

config>router>mpls>mpls-tp

Platforms

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

Description

This command configures the MPLS-TP global ID for the node. The MPLS-TP LSPs originating at this node use this ID as the 'from' global ID. If the *global-id* value is not configured, a value of zero is used.

If an operator expects that inter domain LSPs will be configured, Nokia recommends that the global ID should be set to the local ASN of the node, as configured under **config>system**. If two-byte ASNs are used, the most significant two bytes of the global ID are padded with zeros.

To change the *global-id* value, the **config>router>mpls>mpls-tp** CLI command must be in the **shutdown** state. This state brings down all of the MPLS-TP LSPs on the node. New values are propagated to the system when a **no shutdown** is performed.

Default

no global-id

Parameters

global-id

Specifies the global ID for the node.

Values 0 to 4294967295

node-id

Syntax

node-id node-id

no node-id

Context

config>router>mpls>mpls-tp

Platforms

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

Description

This command configures the MPLS-TP node ID. The MPLS-TP LSPs originating at this node use this ID as the 'from' node ID. The default value of the node ID is the system interface IPv4 address. The node ID may be entered in the 4-octet IPv4 address format, <a.b.c.d>, or as an unsigned 32-bit integer.



Note:

The node ID is not treated as a routable IP address from the perspective of IP routing, and is not advertised in any IP routing protocols.

The MPLS-TP context cannot be administratively enabled unless at least a system interface IPv4 address is configured because MPLS requires that this value is configured.

Default

no node-id

Parameters

node-id

Specifies the MPLS-TP node ID for the node.

Values <a.b.c.d> or 1 to 4294967295

Default System interface IPv4 address

transit-path

Syntax

transit-path *path-name*

no transit-path

Context

config>router>mpls>mpls-tp

Platforms

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

Description

This command enables the configuration or editing of an MPLS-TP transit path at an LSR.

Default

no transit-path

Parameters

path-name

Specifies the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

path-id

Syntax

path-id {*lsp-num* *lsp-num* | **working-path** | **protect-path** [*src-global-id* *src-global-id*] **src-node-id** *src-node-id* **src-tunnel-num** *src-tunnel-num* [*dest-global-id* *dest-global-id*] **dest-node-id** *dest-node-id* [*dest-tunnel-num* *dest-tunnel-num*]}

no path-id

Context

config>router>mpls>mpls-tp>transit-path

Platforms

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

Description

This command configures the path ID for an MPLS-TP transit path at an LSR. The path ID is equivalent to the MPLS-TP LSP ID and is used to generate the maintenance entity group intermediate point (MIP) identifier for the LSP at the LSR. A path ID must be configured for on-demand OAM to verify an LSP at the LSR.

The path ID must contain at least the following parameters: *lsp-num*, *src-node-id*, *src-global-id*, *src-tunnel-num*, and *dest-node-id*.

The path ID must be unique on a node. Nokia recommends that this the configured value is also globally unique.

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

Default

no path-id

Parameters

lsp-num

Specifies the LSP number.

Values 1 to 65535, or working path, or protect-path. A working-path is equivalent to a *lsp-num* of 1, and a protect-path is an *lsp-num* of 2.

src-global-id

Specifies the source global ID.

Values 0 to 4294967295

src-node-id

Specifies the source node ID.

Values a.b.c.d or 1 to 4294967295

src-tunnel-num

Specifies the source tunnel number.

Values 1 to 61440

dest-global-id

Specifies the destination global ID. If the destination global ID is not entered, it is set to the same value as the source global ID.

Values 0 to 4294967295

dest-node-id

Specifies the destination node ID.

Values a.b.c.d or 1 to 4294967295

dest-tunnel-num

Specifies the destination tunnel number. If the destination tunnel number is not entered, it is set to the same value as the source tunnel number.

Values 1 to 61440

forward-path

Syntax

[no] forward-path

Context

config>router>mpls>mpls-tp>transit-path

Platforms

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

Description

This command enables the forward path of an MPLS-TP transit path to be created or edited.

The forward path must be created before the reverse path.

The **no** form of this command removes the forward path. The forward path cannot be removed if a reverse exists.

Default

no forward-path

reverse-path

Syntax

[no] reverse-path

Context

config>router>mpls>mpls-tp>transit-path

Platforms

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

Description

This command enables the reverse path of an MPLS-TP reverse path to be configured or edited.

The reverse path must be created after the forward path. The reverse path must be removed before the forward path.

The **no** form of this command removes the reverse path.

Default

no reverse-path

in-label

Syntax

in-label *in-label* out-label *out-label* out-link *interface-name* [next-hop *next-hop*]

no in-label

Context

config>router>mpls>mpls-tp>transit-path>forward-path

config>router>mpls>mpls-tp>transit-path>reverse-path

Platforms

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

Description

This command configures the label mapping associated with a forward path or reverse path of an MPLS-TP transit path to be configured.

The incoming label, outgoing label, and outgoing interface must be configured using the **in-label**, **out-label**, and **out-link** parameters. If the **out-link** refers to a numbered IP interface, the user may optionally configure the **next-hop** parameter and the system will determine the interface to use to reach the configured next hop, but will check that the user-entered value for the **out-link** corresponds to the link returned by the system. If they do not correspond, the path will not come up.

Default

no in-label

Parameters

in-label

Specifies the incoming label.

Values 32 to 16415

out-label

Specifies the outgoing label.

Values 32 to 16415

interface-name

Specifies the name of the outgoing interface, up to 32 characters, used for the path.

next-hop

Specifies the next hop.

Values a.b.c.d

shutdown

Syntax

[no] shutdown

Context

config>router>mpls>mpls-tp>transit-path

Platforms

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

Description

This command administratively enables or disables an MPLS-TP transit path.

Default

no shutdown

2.17.2.1.6 LSP commands

`lsp`

Syntax

`[no] lsp lsp-name [bypass-only | mpls-tp src-tunnel-num]`

Context

`config>router>mpls`

Platforms

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

Description

This command creates an LSP that is signaled dynamically by the 7210 SAS.

When the LSP is created, the egress router must be specified using the **to** command and at least one primary or secondary path must be specified. All other statements under the LSP hierarchy are optional. The maximum number of static configurable LSPs is 4.

LSPs are created in the administratively down (**shutdown**) state.

The **no** form of this command deletes the LSP. All configuration information associated with this LSP is lost. The LSP must be administratively **shutdown** before it can be deleted.

Parameters

lsp-name

Specifies the name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique.

bypass-only

Keyword to define an LSP as a manual bypass LSP exclusively. When a path message for a new LSP requests bypass protection, the PLR first checks if a manual bypass tunnel satisfying the path constraints exists. If one is found, the 7210 SAS selects it. By default, if no manual bypass tunnel is found, the 7210 SAS dynamically signals a bypass LSP. The CLI for this feature includes a knob that provides the user with the option to disable dynamic bypass creation on a per-node basis.

mpls-tp

Keyword to define an LSP as an MPLS-TP LSP. The following parameters can only be used with an MPLS-TP LSP: **to**, **dest-global-id**, **dest-tunnel-number**, **working-tp-path**, **protect-tp-path**. Other parameters defined for the above LSP types cannot be used. This is supported on 7210 SAS-T network mode, 7210 SAS-R6 and 7210 SAS-R12 devices only.

src-tunnel-num

Specifies the source tunnel number. This is a mandatory time parameter for MPLS-TP LSPs, and has to be assigned by the user based on the configured range of tunnel IDs.

Values 1 to 61440

adaptive

Syntax

[no] adaptive

Context

config>router>mpls>lsp

Platforms

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

Description

This command enables the make-before-break functionality for an LSP or LSP path. When enabled for the LSP, make-before-break is performed for the primary path and all secondary paths of the LSP.

Default

adaptive

adspec

Syntax

[no] adspec

Context

config>router>mpls>lsp

Platforms

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

Description

When enabled, the ADSPEC object will be included in RSVP messages for this LSP. The ADSPEC object is used by the ingress LER to discover the minimum value of the MTU for links in the path of the LSP. By default, the ingress LER derives the LSP MTU from that of the outgoing interface of the LSP path.

A bypass LSP always signals the ADSPEC object because it protects primary paths that signal the ADSPEC object and primary paths that do not. This means that MTU of the LSP at ingress LER may change to a different value from that derived from the outgoing interface even if the primary path has ADSPEC disabled.

Default

no adspec

bgp-transport-tunnel

Syntax

bgp-transport-tunnel include | exclude

Context

config>router>mpls>lsp

Platforms

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

Description

This command controls whether RSVP-TE LSP can be used as a transport LSP for BGP tunnel routes.

Default

bgp-transport-tunnel exclude

Parameters

include

Keyword to enable RSVP-TE LSP to be used as transport LSP from the ASBR to local PE router, from ingress PE to ASBR in the local AS, or between multi-hop eBGP peers with ASBR to ASBR adjacency.

exclude

Disables RSVP-TE LSP from being used as transport LSP from the ASBR to local PE router, from ingress PE to ASBR in the local AS, or between multi-hop eBGP peers with ASBR to ASBR adjacency.

cspf

Syntax

[no] cspf [use-te-metric]

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Platforms

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

Description



Note:

- The **config>router>mpls>lsp-template** context is only supported on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T (operating in network mode), and 7210 SAS-Sx/S 1/10GE (operating in standalone and standalone-VC mode).
- In the **lsp-template** context, this command is only supported with NG-MVPN. It is not supported with other applications.

This command enables constrained shortest path first (CSPF) computation for constrained-path LSPs. Constrained-path LSPs take configuration constraints into account. CSPF is also used to calculate the detour routes when the **fast-reroute** command is enabled.

Explicitly configured LSPs for which each hop from ingress to egress is specified do not use CSPF. The LSP is set up using RSVP signaling from ingress to egress.

If an LSP is configured with the **fast-reroute** *frr-method* option specified but does not enable CSPF, neither global revertive nor local revertive will be available for the LSP to recover.

The **no** form of this command disables CSPF computation for constrained-path LSPs.

Default

no cspf

Parameters

use-te-metric

Keyword to use the TE metric for the CSPF computation of the LSP path.

exclude

Syntax

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

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Platforms

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

Description



Note:

- The **config>router>mpls>lsp-template** context is only supported on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T (operating in network mode), and 7210 SAS-Sx/S 1/10GE (operating in standalone and standalone-VC mode).
- In the **lsp-template** context, this command is only supported with NG-MVPN. It is not supported with other applications.

This command specifies the admin groups to be excluded when an LSP is set up in the primary or secondary contexts. A maximum of 5 groups can be specified per single operation of the **exclude** command. However, a maximum of 32 groups can be specified per LSP through multiple operations. The admin groups are defined in the **config>router>if-attribute** context.

The **no** form of this command removes the admin groups configured for exclusion.

Default

no exclude

Parameters

group-name

Specifies the existing group name, up to 32 characters, to be excluded when an LSP is set up.

fast-reroute

Syntax

fast-reroute [*frr-method*]

no fast-reroute

Context

config>router>mpls>lsp

Platforms

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

Description

This command enables a precomputed detour LSP from each node in the path of the LSP. In case of failure of a link or LSP between two nodes, traffic is immediately rerouted on the precomputed detour LSP, which avoids packet loss.

When the **fast-reroute** command is enabled, each node along the path of the LSP tries to establish a detour LSP as follows.

- Each upstream node sets up a detour LSP that avoids only the immediate downstream node, and merges back on to the main path of the LSP as soon as possible.

If it is not possible to set up a detour LSP that avoids the immediate downstream node, a detour can be set up to the downstream node on a different interface.

- The detour LSP may take one or more hops (see [hop-limit](#)) before merging back on to the main LSP path.
- When the upstream node detects a downstream link or node failure, the ingress router switches traffic to a standby path if one was set up for the LSP.

Fast reroute is available only for the primary path. No configuration is required on the transit hops of the LSP. The ingress router will signal all intermediate routers using RSVP to set up their detours. TE must be enabled for fast-reroute to work.

If an LSP is configured with the **fast-reroute** *frr-method* option specified but does not enable CSPF, neither global revertive nor local revertive will be available for the LSP to recover.

The **no** form of this command removes the detour LSP from each node on the primary path. This command will also remove configuration information about the hop-limit and the bandwidth for the detour routes.

The **no** form of **fast-reroute hop-limit** command reverts to the default value.

Default

no fast-reroute

Parameters

frr-method

Specifies the fast reroute method.

Values **one-to-one** — Keyword to specify that a label-switched path is established that intersects the original LSP somewhere downstream of the point of the link or node failure. For each LSP that is backed up, a separate backup LSP is established.

facility — Keyword, sometimes called many-to-one, that takes advantage of the MPLS label stack. Instead of creating a separate LSP for every backed-up LSP, a single LSP is created and serves to back up a set of LSPs. This LSP tunnel is called a bypass tunnel.

hop-limit

Syntax

hop-limit *limit*

no hop-limit

Context

config>router>mpls>lsp>fast-reroute

Platforms

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

Description

This command configures the fast reroute context to set how many more routers a detour is allowed to traverse compared to the LSP itself. For example, if an LSP traverses four routers, any detour for the LSP can be no more than ten router hops, including the ingress and egress routers.

Default

hop-limit 16

Parameters

limit

Specifies the maximum number of hops.

Values 0 to 255

node-protect

Syntax

[no] node-protect

Context

config>router>mpls>lsp>fast-reroute

Platforms

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

Description

This command enables node and link protection on the specified LSP. Node protection ensures that traffic from an LSP traversing a neighboring router will reach its destination even if the neighboring router fails.

The **no** form of this command disables node and link protection on the specified LSP.

Default

node-protect

propagate-admin-group

Syntax

[no] propagate-admin-group

Context

config>router>mpls>lsp>fast-reroute

Platforms

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

Description

The command enables the signaling of the primary LSP path admin-group constraints in the FRR object at ingress.

When this command is executed, the admin-group constraints configured in the context of the P2P LSP primary path, or the constraints configured in the context of the LSP and inherited by the primary path, are copied into the FAST_REROUTE object. The admin-group constraints are copied into the "include-any" or "exclude-any" fields.

During LSP signaling to the downstream node, the ingress LER also propagates the admin-group constraints, which allows the node to include these constraints in the selection of the FRR backup LSP for LSP primary path protection.

The ingress LER inserts the FAST_REROUTE object, by default, in a primary LSP path message. If the user disables the object using **config>router>mpls>no frr-object** command, the admin-group constraints are not propagated.

The same admin-group constraints can be copied into the Session Attribute object for use by an LSR, typically an ABR, to expand the ERO of an inter-area LSP path. The constraints are also used by any LSR node in the path of a CSPF or non-CSPF LSP to check the admin-group constraints against the ERO regardless if the hop is strict or loose. These constraints are governed strictly by the **config>router>mpls>lsp>propagate-admin-group** command.

That is, the user can copy the primary path admin-group constraints into only the FAST_REROUTE object only, or only the Session Attribute object only, or both. However, the PLR rules for processing the admin-group constraints can make use of either of the two object admin-group constraints.

This feature is supported with the primary path of an RSVP P2P LSP in both intra-area and inter-area TE.

The **no** form of this command disables administrative group constraint signaling in the FRR object.

Default

no propagate-admin-group

from

Syntax

from *ip-address*

Context

config>router>mpls>lsp

Platforms

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

Description

This optional command configures the IP address of the ingress router for the LSP. When this command is not specified, the system IP address is used. IP addresses that are not defined in the system are allowed. If an invalid IP address is entered, LSP bring-up fails and an error is logged.

If an interface IP address is specified as the **from** address, and the egress interface of the next-hop IP address is a different interface, the LSP is not signaled. As the egress interface changes because of changes in the routing topology, an LSP recovers if the **from** IP address is the system IP address and not a specific interface IP address.

Only one **from** address can be configured.

Parameters

ip-address

Specifies the IP address of the ingress router. This can be either the interface or the system IP address. If the IP address is local, the LSP must egress through that local interface which ensures local strictness.

Values System IP or network interface IP addresses

Default System IP address

hop-limit

Syntax

hop-limit *number*

no hop-limit

Context

config>router>mpls>lsp

config>router>mpls>lsp>fast-reroute

Platforms

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

Description

This command configures the maximum number of hops that an LSP can traverse, including the ingress and egress routers. An LSP is not set up if the hop limit is exceeded. This value can be changed dynamically for an LSP that is already set up with the following implications.

If the new value is less than the current number of hops of the established LSP, the LSP is brought down. Software then tries to re-establish the LSP within the new **hop-limit** *number*. If the new value is equal to or greater than the current number hops of the established LSP, the LSP is not affected.

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

Default

hop-limit 255

Parameters

number

Specifies the number of hops the LSP can traverse, expressed as an integer.

Values 2 to 255

include

Syntax

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

Context

config>router>mpls>lsp

config>router>mpls>lsp>primary

config>router>mpls>lsp>secondary

config>router>mpls>lsp-template

Platforms

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

Description



Note:

- The **config>router>mpls>lsp-template** context is only supported on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T (operating in network mode), and 7210 SAS-Sx/S 1/10GE (operating in standalone and standalone-VC mode).
- In the **lsp-template** context, this command is only supported with NG-MVPN. It is not supported with other applications.

This command configures the admin groups to be included when an LSP is set up. Up to 5 groups per operation can be specified, up to 32 maximum.

The **no** form of this command deletes the specified groups in the specified context.

Default

no include

Parameters

group-name

Specifies admin groups, up to 32 characters, to be included when an LSP is set up.

ldp-over-rsvp

Syntax

[no] ldp-over-rsvp [include | exclude]

Context

config>router>mpls>lsp

Platforms

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

Description

This command configures if this LSP will be included in LDP over RSVP.

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

Default

no ldp-over-rsvp

Parameters

include

Specifies that this LSP will be included in LDP over RSVP.

exclude

Specifies that this LSP will be excluded from LDP over RSVP.

metric

Syntax

metric *metric*

Context

config>router>mpls>lsp

Platforms

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

Description

This command configures the metric for this LSP, which is used to select an LSP among a set of LSPs that are destined for the same egress router. The LSP with the lowest metric is selected.

In LDP-over-RSVP, LDP performs a lookup in the Routing Table Manager (RTM), which provides the next hop to the destination PE and the advertising router (ABR or destination PE). If the advertising router matches the targeted LDP peer, LDP performs a second lookup for the advertising router in the Tunnel

Table Manager (TTM). This lookup returns the best RSVP LSP to use to forward packets for an LDP FEC learned through the targeted LDP session. The lookup returns the LSP with the lowest metric. If multiple LSPs have the same metric, the result of the lookup is to select the first available LSP in the TTM.

Default

metric 1

Parameters

metric

Specifies the metric for this LSP, which is used to select an LSP among a set of LSPs that are destined for the same egress router.

Values 1 to 65535

path-profile

Syntax

path-profile *profile-id* [**path-group** *group-id*]

no path-profile *profile-id*

Context

config>router>mpls>lsp

Platforms

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

Description

This command configures the PCE path profile and path group ID.

The PCE supports the computation of disjoint paths for two LSPs originating or terminating on the same or different PE routers. To indicate this constraint to the PCE, the user configures the PCE path profile ID and path group ID to which the PCE-computed or PCE-controlled LSP belongs. Because the PCC passes these parameters transparently to the PCE, the parameters are opaque data to the router.

The association of the optional path group ID allows the PCE to determine the profile ID to use with this path group ID. Although one path group ID is allowed per profile ID, you can execute the **path-profile** command multiple times and enter the same path group ID with multiple profile IDs. A maximum of five **path-profile** *profile-id* [**path-group** *group-id*] entries can be associated with the same LSP.

Parameters

profile-id

Specifies the profile ID.

Values 1 to 4294967295

group-id

Specifies the path group ID.

Values 0 to 4294967295

pce-computation

Syntax

[no] pce-computation

Context

config>router>mpls>lsp

Platforms

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

Description

This command enables the PCE-computed LSP mode of operation for an RSVP-TE LSP.

The user can grant only path computation requests (PCE-computed) or both path computation requests and path updates (PCE-controlled) to a PCE for a specific LSP.

The **pce-computation** command sends the path computation request to the PCE instead of the local CSPF. Enabling this option allows the PCE to perform path computations for the LSP at the request of the PCC router only. This feature is used in cases where the operator wants to use the PCE-specific path computation algorithm instead of the local router CSPF algorithm.

The default configuration is **no pce-computation**. To enable the **pce-computation** command or **pce-control** command, you must first enable the **cspf** option, or this configuration is rejected. Conversely, an attempt to disable the **cspf** option on an RSVP-TE LSP that has the **pce-computation** command or **pce-control** command enabled is rejected.

Default

no pce-computation

pce-control

Syntax

[no] pce-control

Context

config>router>mpls>lsp

Platforms

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

Description

This command enables the PCE-controlled LSP mode of operation for an RSVP-TE LSP.

Using the **pce-control** command, the PCC router delegates full control of the LSP to the PCE (PCE-controlled). As a result, PCE acts in an active stateful mode for this LSP. The PCE can reroute the path following a failure or reoptimize the path and update the router without an update request from the PCC router.

The user can delegate CSPF and non-CSPF LSPs, or LSPs that have the **pce-computation** option enabled or disabled. The LSP maintains the latest active path computed by the PCE or the PCC router at the time it is delegated. The PCE will only update the path at the next network event or reoptimization.

The default configuration is **no pce-control**. To enable the **pce-control** command or **pce-computation** command, you must first enable the **cspf** option; otherwise, this configuration is rejected. Conversely, an attempt to disable the **cspf** option on an RSVP-TE LSP that has the **pce-control** command or **pce-computation** command enabled is rejected.

If PCE reporting is disabled for the LSP, either because of inheritance from the MPLS-level configuration or because of LSP-level configuration, enabling the **pce-control** option for the LSP has no effect.

Default

no pce-control

pce-report

Syntax

pce-report {enable | disable | inherit}

Context

config>router>mpls>lsp

Platforms

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

Description

This command configures the reporting mode to a PCE for an RSVP-TE LSP.

The PCC LSP database is synchronized with the PCE LSP database using the PCEP PCRpt (PCE Report) message for PCC-controlled, PCE-computed, and PCE-controlled LSPs.

Use the global MPLS-level **pce-report** command (**config>router>mpls>pce-report**) to enable or disable PCE reporting for all RSVP-TE LSPs during PCE LSP database synchronization.

The LSP-level **pce-report** command overrides the global configuration for reporting an LSP to the PCE. The default configuration is to inherit the global MPLS-level configuration. The **inherit** option reconfigures the LSP to inherit the global configuration.

Default

pce-report inherit

Parameters

enable

Keyword to enable PCE reporting.

disable

Keyword to disable PCE reporting.

inherit

Keyword to inherit the global configuration for PCE reporting.

propagate-admin-group

Syntax

[no] **propagate-admin-group**

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Platforms

7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12 (for **config>router>mpls>lsp** context)

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC mode) (for **config>router>mpls>lsp-template** context)

Description



Note:

In the **lsp-template** context, this command is only supported with NG-MVPN. It is not supported with other applications.

This command enables the propagation of session attribute objects with resource affinity (C-type 1) in a Path message. If a session attribute with resource affinity is received at an LSR, the LSR checks the compatibility of admin-groups received in the Path message with configured admin-groups on the egress interface of the LSP.

To support admin-groups for inter-area LSPs, the ingress node must configure the propagation of admin-groups within the SESSION_ATTRIBUTE object. If a Path message is received by an LSR node that has the **cspf-on-loose-hop** command configured and the message includes admin-groups, the ERO expansion by CSPF to calculate the path to the next loose hop will include the admin-group constraints received from the ingress node.

If the **cspf-on-loose-hop** command is disabled, the SESSION_ATTRIBUTE object without resource affinity (C-Type 7) is propagated in the Path message, and CSPF at the LSR node will not include admin group constraints.

Admin-group propagation is supported with P2P LSPs.

The user can change the value of the **propagate-admin-group** option on the fly. An RSVP P2P LSP performs a make-before-break (MBB) when changing the configuration.

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

Default

no propagate-admin-group

retry-limit

Syntax

retry-limit *number*

no retry-limit

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Platforms

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

Description



Note:

- The **config>router>mpls>lsp-template** context is only supported on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T (operating in network mode), and 7210 SAS-Sx/S 1/10GE (operating in standalone and standalone-VC mode).
- In the **lsp-template** context, this command is only supported with NG-MVPN. It is not supported with other applications.

This command configures the number of attempts the software should make to re-establish the LSP after the LSP has failed. After each successful attempt, the counter is reset to zero.

When the configured retry limit is reached, no more attempts are made and the LSP path is set to the **shutdown** state.

Use the **config>router>mpls>lsp>no shutdown** command to bring up the path after the retry limit is exceeded.

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

Default

retry-limit 0

Parameters

number

Specifies the number of software attempts to re-establish the LSP after it has failed. A value of 0 indicates to retry forever.

Values 0 to 10000

retry-timer

Syntax

retry-timer *seconds*

no retry-timer

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Platforms

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

Description



Note:

- The **config>router>mpls>lsp-template** context is only supported on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T (operating in network mode), and 7210 SAS-Sx/S 1/10GE (operating in standalone and standalone-VC mode).
- In the **lsp-template** context, this command is only supported with NG-MVPN. It is not supported with other applications.

This command configures the time, in seconds, between LSP re-establishment attempts after the LSP has failed.

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

Default

retry-timer 30

Parameters

seconds

Specifies the amount of time, in seconds, between attempts to re-establish the LSP after it has failed.

Values 1 to 600

rsvp-resv-style

Syntax

rsvp-resv-style [se | ff]

Context

config>router>mpls>lsp

Platforms

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

Description

This command configures the RSVP reservation style, shared explicit (se) or fixed filter (ff). A reservation style is a set of control options that specify a number of supported parameters. The style information is part of the LSP configuration.

Default

rsvp-resv-style se

Parameters

ff

Keyword to configure fixed filter reservation style, which is a single reservation with an explicit scope. This reservation style specifies an explicit list of senders and a distinct reservation for each of them. A specific reservation request is created for data packets from a particular sender. The reservation scope is determined by an explicit list of senders.

se

Keyword to configure shared explicit reservation style, which is a shared reservation with a limited scope. This reservation style specifies a shared reservation environment with an explicit reservation scope. This reservation style creates a single reservation over a link that is shared by an explicit list of senders. Because each sender is explicitly listed in the RESV message, different labels can be assigned to different sender-receiver pairs, thereby creating separate LSPs.

shutdown

Syntax

[no] shutdown

Context

config>router>mpls>lsp

Platforms

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

Description

This command disables the existing LSP, including the primary path and any standby secondary paths.

To shut down only the primary path, enter the **config>router>mpls>lsp>primary> shutdown** command.

To shut down a specific standby secondary path, enter the **config>router>mpls>lsp>secondary>shutdown** command. The existing configuration of the LSP is preserved.

Use the **no** form of this command to restart the LSP. LSPs are created in a **shutdown** state. Use this command to administratively bring up the LSP.

Default

shutdown

to

Syntax

to *ip-address*

Context

config>router>mpls>lsp

Platforms

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

Description

This command configures the system IP address of the egress router for the LSP. This command is mandatory to create an LSP.

An IP address for which a route does not exist is allowed in the configuration. If the LSP signaling fails because the destination is not reachable, an error is logged and the LSP operational status is set to down.

Parameters

ip-address

Specifies the system IP address of the egress router.

vprn-auto-bind

Syntax

vprn-auto-bind [include | exclude]

Context

config>router>mpls>lsp

Platforms

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

Description

This command configures whether the associated LSP can be used as part of the auto-bind feature for VPRN services. By default, a named LSP is allowed to be used for the auto-bind feature.

When the **vprn-auto-bind** command is set to **exclude**, the associated LSP is not used by the auto-bind feature within VPRN services.

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

Default

vprn-auto-bind include

Parameters

include

Keyword to allow an associated LSP to be used by auto-bind for VPRN services.

exclude

Keyword to prevent the associated LSP from being used with the auto-bind feature for VPRN services.

lsp-template

Syntax

lsp-template *template-name* **p2mp**

no lsp-template *template-name*

Context

config>router>mpls

Platforms

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

Description

This command creates a template that can be referenced by a client application where dynamic LSP creation is required. The LSP template type **p2mp** is mandatory.



Note:

The **lsp-template** command is only supported with NG-MVPN. This command is not supported for other applications.

The **no** form of this command deletes the LSP template. An LSP template cannot be deleted if a client application is using it.

Parameters

template-name

Specifies the name of the LSP template, up to 32 characters. An LSP template name and LSP name must not be the same.

p2mp

Mandatory keyword to configure P2MP as the LSP type that this template will signal.

default-path

Syntax

default-path *path-name*

Context

config>router>mpls>lsp-template

Platforms

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

Description

A default path binding must be provided before the LSP template can be used for signaling LSP. The LSP template must be shut down to modify **default-path** binding.



Note:

In the **lsp-template** context, this command is only supported with NG-MVPN. It is not supported with other applications.

Parameters

path-name

Specifies the default path binding, up to 32 characters.

2.17.2.1.7 Primary and secondary path commands

primary

Syntax

primary *path-name*

no primary

Context

config>router>mpls>lsp

Platforms

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

Description

This command configures a preferred path for the LSP. This command is optional only if the secondary path name is included in the LSP definition. Only one primary path can be defined for an LSP.

Some of the attributes of the LSP, such as the bandwidth and hop limit, can be optionally specified as the attributes of the primary path. The attributes specified in the **primary path-name** command override the LSP attributes.

The **no** form of this command deletes the association of this *path-name* from the **lsp lsp-name**. All configurations specific to this primary path, such as record, bandwidth, and hop limit, are deleted. The primary path must be shut down to delete it. The **no primary** command will not result in any action except a warning message on the console indicating that the primary path is administratively up.

Parameters

path-name

Specifies the case-sensitive alphanumeric name label for the LSP path, up to 32 characters in length.

secondary

Syntax

[no] **secondary path-name**

Context

config>router>mpls>lsp

Platforms

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

Description

This command configures an alternative path that the LSP uses if the primary path is not available. This command is optional and is not required if the **config router mpls lsp lsp-name primary path-name** command is specified. After the switch over from the primary to the secondary path, the software continuously tries to revert to the primary path. The switch back to the primary path is based on the **retry-timer** interval.

Up to eight secondary paths can be specified. All the secondary paths are considered equal and the first available path is used. The software will not switch back among secondary paths.

Software starts the signaling of all non-standby secondary paths at the same time. Retry counters are maintained for each unsuccessful attempt. Once the retry limit is reached on a path, software will not attempt to signal the path and administratively shuts down the path. The first successfully established path is made the active path for the LSP.

The **no** form of this command removes the association between this path-name and lsp-name. All specific configurations for this association are deleted. The secondary path must be shut down first to delete it. The **no secondary path-name** command will not result in any action except a warning message on the console indicating that the secondary path is administratively up.

Parameters

path-name

Specifies the case-sensitive alphanumeric name label for the LSP path, up to 32 characters in length.

adaptive

Syntax

[no] adaptive

Context

config>router>mpls>lsp>primary

config>router>mpls>lsp>secondary

Platforms

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

Description

This command enables the make-before-break functionality for an LSP or a primary or secondary LSP path. When enabled for the LSP, make-before-break is performed for the primary path and all the secondary paths of the LSP.

Default

adaptive

working-tp-path

Syntax

[no] working-tp-path

Context

config>router>mpls>lsp

Platforms

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

Description

This command configures or edits the working path for an MPLS-TP LSP. At least one working path (but not more than one working path) must be created for an MPLS-TP LSP. If MPLS-TP linear protection is also configured, this is the path that is used as the default working path for the LSP, and it must be created before the protect path. The **working-tp-path** can only be deleted if no **protect-tp-path** exists for the LSP.

The following commands are applicable to the **working-tp-path**: **lsp-num**, **in-label**, **out-label**, **mep**, **shutdown**.

Default

no working-tp-path

protect-tp-path

Syntax

[no] protect-tp-path

Context

config>router>mpls>lsp

Platforms

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

Description

This command configures or edits the protect path for an MPLS-TP LSP. At least one working path must exist before a protect path can be created for an MPLS-TP LSP. If MPLS-TP linear protection is also configured, this is the path that is used as the default protect path for the LSP. The protect path must be deleted before the working path. Only one protect path can be created for each MPLS-TP LSP.

The following commands are applicable to the **working-tp-path**: **lsp-num**, **in-label**, **out-label**, **mep**, and **shutdown**.

lsp-num

Syntax

lsp-num **lsp-num**

no **lsp-num**

Context

config>router>mpls>lsp>working-tp-path

config>router>mpls>lsp>protect-tp-path

Platforms

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

Description

This command configures the MPLS-TP LSP number for the working TP path or the protect TP path.

Default

no lsp-num

Parameters

lsp-num

Specifies the LSP number.

Values 1 to 65535

Default 1 for a working path, 2 for a protect path

in-label

Syntax

in-label in-label

Context

config>router>mpls>lsp>working-tp-path
config>router>mpls>lsp>protect-tp-path

Platforms

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

Description

This command configures the incoming label for the reverse path, working path, or protect path of an MPLS-TP LSP. MPLS-TP LSPs are bidirectional, and so an incoming label value must be specified for each path.

Default

no in-label

Parameters

in-label

Specifies the in label.

Values 32 to 16415

out-label

Syntax

out-label out-label out-link if-name [next-hop ip-address]
no out-label

Context

config>router>mpls>lsp>working-tp-path
config>router>mpls>lsp>protect-tp-path

Platforms

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

Description

This command configures the outgoing label value to use for an MPLS-TP working or protect path. The **out-link** is the outgoing interface on the node that this path will use, and must be specified. If the **out-link** refers to a numbered IP interface, the user may optionally configure the **next-hop** parameter and the system will determine the interface to use to reach the configured **next-hop**, but will check that the user-entered value for the **out-link** corresponds to the link returned by the system. If they do not correspond, the path will not come up.

Default

no out-label

Parameters

out-label

Specifies the out label.

Values 32 to 16415

if-name

Specifies the interface name.

ip-address

Specifies the IPv4 address in a.b.c.d.

mep

Syntax

[no] mep

Context

config>router>mpls>lsp>working-tp-path

config>router>mpls>lsp>protect-tp-path

Platforms

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

Description

This command creates or edits an MPLS-TP maintenance entity group (MEG) endpoint (MEP) on an MPLS-TP path. MEPs represent the termination point for OAM flowing on the path, as well as linear protection for the LSP. Only one MEP can be configured at each end of the path.

The following commands are applicable to a MEP on an MPLS-TP working or protect path: **oam-template**, **bfd-enable**, and **shutdown**. In addition, a protection-template may be configured on a protect path.

The **no** form of this command removes a MEP from an MPLS-TP path.

oam-template

Syntax

oam-template *name*

no oam-template

Context

config>router>mpls>lsp>working-tp-path

config>router>mpls>lsp>protect-tp-path

Platforms

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

Description

This command applies an OAM template to an MPLS-TP working or protect path. It contains configuration parameters for proactive OAM mechanisms that can be enabled on the path, for example, BFD. Configuration of an OAM template is optional.

The **no** form of this command removes the OAM template from the path.

Default

no oam-template

Parameters

name

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

shutdown

Syntax

[no] shutdown

Context

config>router>mpls>lsp>working-tp-path>mep

config>router>mpls>lsp>protect-tp-path>mep config>router>mpls>lsp>working-tp-path

config>router>mpls>lsp>protect-tp-path

Platforms

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

Description

This command disables the existing LSP, including the primary path and any standby secondary paths.

To shut down only the primary path, enter the **config router mpls lsp *lsp-name* primary *path-name* shutdown** command.

To shut down a specific standby secondary path, enter the **config router mpls lsp *lsp-name* secondary *path-name* shutdown** command. The existing configuration of the LSP is preserved.

The **no** form of this command restarts the LSP. LSPs are created in a **shutdown** state. Use this command to administratively bring up the LSP.

Default

shutdown

bfd-enable

Syntax

bfd-enable *bfd-mode*

no bfd-enable

Context

config>router>mpls>lsp>working-tp-path>mep

config>router>mpls>lsp>protect-tp-path>mep

Platforms

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

Description

This command associates the operational state of an MPLS-TP path with a BFD session for which control packets flow on the path. The BFD packets are encapsulated in a generic associated channel (G-ACh) on the path. The timer parameters of the BFD session are taken from the OAM template of the MEP.

A value of **cc** means that the BFD session is only used for continuity check of the MPLS-TP path. In this case, the **cc** timer parameters of the OAM template apply. A value of **cc_cv** means that the BFD session is used for both continuity checking and connectivity verification, and the **cc_cv** timers of the OAM template apply.

This form of this **bfd-enable** command is only applicable when it is configured under a MEP used on an MPLS-TP working or protect path.

Default

no bfd-enable

Parameters

bfd-mode

Specifies the BFD mode.

Values **cc** — Option to indicate that BFD runs in CC only mode. This mode uses GACH channel type 0x07.

cc_cv — Option to indicate that BFD runs in combined CC and CV mode. This mode uses channel type 0x22 for MPLS-TP CC packets, and 0x23 for MPLS-TP CV packets.

protection-template

Syntax

protection-template name
no protection-template

Context

config>mpls>lsp>protect-tp-path

Platforms

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

Description

This command applies a protection template name to an MPLS-TP LSP under which the protect path is configured. If the template is applied, MPLS-TP 1:1 linear protection is enabled on the LSP using the parameters specified in the named template.

A named protection template can only be applied to the protect path context of an MPLS-TP LSP.

The **no** form of this command removes the template and disables MPLS-TP linear protection on the LSP.

Default

no protection-template

Parameters

name

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

bandwidth

Syntax

bandwidth rate-in-mbps
no bandwidth

Context

config>router>mpls>lsp>primary

```
config>router>mpls>lsp>secondary
```

Platforms

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

Description

This command specifies the amount of bandwidth to be reserved for the LSP path.

The **no** form of this command resets bandwidth parameters (no bandwidth is reserved). This is the bandwidth setting in the global LSP configuration.

Default

no bandwidth

Parameters

rate-in-mbps

Specifies the amount of bandwidth reserved for the LSP path in Mbps.

Values 0 to 100000

exclude

Syntax

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

Context

```
config>router>mpls>lsp>primary
```

```
config>router>mpls>lsp>secondary
```

Platforms

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

Description

This command specifies the admin groups to be excluded when an LSP is set up. Up to five groups per operation can be specified, up to 32 maximum. The admin groups are defined in the **config>router>if-attribute** context.

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

Default

no exclude

Parameters

group-name

Specifies the existing group name to be excluded when an LSP is set up.

hop-limit

Syntax

hop-limit *number*

no hop-limit

Context

config>router>mpls>lsp>primary

config>router>mpls>lsp>secondary

Platforms

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

Description

This optional command overrides the **config router mpls lsp /sp-name hop-limit** command. This command specifies the total number of hops that an LSP traverses, including the ingress and egress routers.

This value can be changed dynamically for an LSP that is already set up with the following implications.

If the new value is less than the current hops of the established LSP, the LSP is brought down. MPLS then tries to re-establish the LSP within the new hop limit number. If the new value is equal to or more than the current hops of the established LSP, the LSP will be unaffected.

The **no** form of this command reverts to the default values defined using the **config router mpls lsp /sp-name hop-limit** command.

Default

no hop-limit

Parameters

number

Specifies the number of hops the LSP can traverse, expressed as an integer.

Values 2 to 255

path-preference

Syntax

[no] **path-preference** *value*

Context

config>router>mpls>lsp>secondary

Platforms

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

Description

This command enables the use of path preference among configured standby secondary paths for each LSP. If all standby secondary paths have a default path preference value, a non-standby secondary path remains an active path, while a standby secondary is available. A standby secondary path configured with highest priority (lowest path preference value) must be made the active path when the primary path is not in use. Path preference can be configured on standby secondary path.

The **no** form of this command resets the path preference to the default value.

Default

path-preference 255

Parameters

value

Specifies an alternate path for the LSP if the primary path is not available.

Values 1 to 255

record

Syntax

[no] record

Context

config>router>mpls>lsp>primary

config>router>mpls>lsp>secondary

config>router>mpls>lsp-template

Platforms

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

Description



Note:

- The **config>router>mpls>lsp-template** context is only supported on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T (operating in network mode), and 7210 SAS-Sx/S 1/10GE (operating in standalone and standalone-VC mode).
- In the **lsp-template** context, this command is only supported with NG-MVPN. It is not supported with other applications.

This command enables recording of all hops that an LSP path traverses. Enabling **record** increases the size of the PATH and RESV refresh messages for the LSP because this information is carried end-to-end along the LSP path. The increase in control traffic for each LSP may impact scalability.

The **no** form of this command disables the recording of all hops for a specific LSP. There are no restrictions for the **no** command usage.

The **no** form of this command also disables the **record-label** command.

Default

record

record-label

Syntax

[no] record-label

Context

config>router>mpls>lsp>primary

config>router>mpls>lsp>secondary

config>router>mpls>lsp-template

Platforms

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

Description



Note:

- The **config>router>mpls>lsp-template** context is only supported on the 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-T (operating in network mode), and 7210 SAS-Sx/S 1/10GE (operating in standalone and standalone-VC mode).
- In the **lsp-template** context, this command is only supported with NG-MVPN. It is not supported with other applications.

This command enables recording of all labels at each node that an LSP path traverses. Enabling the **record-label** command also enables the **record** command if it is not already enabled.

The **no** form of this command disables the recording of hops that an LSP path traverses.

Default

record-label

srlg

Syntax

[no] srlg

Context

config>router>mpls>lsp>secondary

Platforms

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

Description

This command enables the use of the SRLG constraint in the computation of a secondary path for an LSP at the head-end LER. When this feature is enabled, CSPF includes the SRLG constraint in the computation of the secondary LSP path.

CSPF requires that the primary LSP be established already and in the up state, because the head-end LER needs the most current ERO computed by CSPF for the primary path and CSPF includes the list of SRLGs in the ERO during the CSPF computation of the primary path. At a subsequent establishment of a secondary path with the SRLG constraint, the MPLS/RSVP task queries CSPF again, which provides the list of SRLG group numbers to be avoided. CSPF prunes all links with interfaces that belong to the same SRLGs as the interfaces included in the ERO of the primary path. If CSPF finds a path, the secondary is setup. If CSPF does not find a path, MPLS/RSVP keeps retrying the requests to CSPF.

If CSPF is not enabled on the LSP (using the **lsp *lsp-name* cspf** command), a secondary path of that LSP that includes the SRLG constraint is shut down and a specific failure code indicates the exact reason for the failure in the **show>router>mpls>lsp path detail** output.

At initial primary LSP path establishment, if primary does not come up or is not configured, the SRLG secondary is not signaled and is put in the down state. A specific failure code indicates the exact reason for the failure in the **show>router>mpls>lsp path detail** output. However, if a non-SRLG secondary path was configured, such as a secondary path with the SRLG option disabled, the MPLS/RSVP task signals it and the LSP uses it.

As soon as the primary path is configured and successfully established, MPLS/RSVP moves the LSP to the primary path and signals all SRLG secondary paths.

Any time the primary path is reoptimized, has undergone MBB operation, or has come back up after being down, the MPLS/RSVP task checks with CSPF to determine if the SRLG secondary path should be resignaled. If the MPLS/RSVP task finds that the current secondary path is no longer SRLG disjoint — for example, the path became ineligible — it puts the path on a delayed MBB immediately after the expiry of the retry timer. If MBB fails on the first try, the secondary path is torn down and the path is put on retry.

At the next opportunity that the primary goes down, the LSP uses an eligible SRLG secondary path if the path is in the up state. If all secondary eligible SRLG paths are in the down state, MPLS/RSVP uses a non-SRLG secondary path if the path is configured and in the up state. If, while the LSP is using a non-SRLG secondary path, an eligible SRLG secondary path comes back up, MPLS/RSVP will not switch the path of the LSP to it. As soon as the primary path is resignaled and comes up with a new SRLG list, MPLS/RSVP resignals the secondary path using the new SRLG list.

A secondary path that becomes ineligible as a result of an update to the SRLG membership list of the primary path will have the ineligibility status removed when any of the following events occur.

- A successful MBB operation of the standby SRLG path occurs, making the path eligible again.
- The standby path goes down, in which case MPLS/RSVP puts the standby on retry when the retry timer expires. If successful, it becomes eligible. If not successful after the retry-timer expires or the number of retries reaches the number configured under the retry-limit parameter, it is left down.
- The primary path goes down, in which case the ineligible secondary path is immediately torn down and will only be resigaled when the primary path comes back up with a new SRLG list.

After the primary path of the LSP is set up and is operationally up, any subsequent changes to the SRLG group membership of an interface that the primary path is using is not considered until the next opportunity that the primary path is resigaled. The primary path may be resigaled because of a failure or to a make-before-break operation. A make-before-break operation occurs as a result of a global revertive operation, a timer-based or manual re-optimization of the LSP path, or a change by a user to any of the path constraints.

After an SRLG secondary path is setup and is operationally up, any subsequent changes to the SRLG group membership of an interface that the secondary path is using is not considered until the next opportunity that the secondary path is resigaled. The secondary path is resigaled because of a failure, to a resignaling of the primary path, or to a make-before-break operation. A make-before break operation occurs as a result of a timer-based or manual reoptimization of the secondary path, or an operator change to any of the path constraints of the secondary path, including enabling or disabling the SRLG constraint itself.

In addition, the user-configured include or exclude admin group statements for this secondary path are also checked along with the SRLG constraints by CSPF.

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

Default

no srlg

standby

Syntax

[no] standby

Context

config>router>mpls>lsp>secondary

Platforms

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

Description

The secondary path LSP is normally signaled when the primary path LSP fails. The **standby** keyword ensures that the secondary path LSP is signaled and maintained indefinitely in a hot-standby state. When the primary path is re-established, the traffic is switched back to the primary path LSP.

The **no** form of this command specifies that the secondary LSP is signaled when the primary path LSP fails.

2.17.2.1.8 LSP path commands

hop

Syntax

hop *hop-index ip-address* {**strict** | **loose**}

no hop *hop-index*

Context

config>router>mpls>path

Platforms

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

Description

This command configures the IP address of the hops that the LSP should traverse on its way to the egress router. The IP address can be the interface IP address or the system IP address. If the system IP address is specified, the LSP can choose the best available interface.

Optionally, the LSP ingress and egress IP address can be included as the first and last hop. A hop list can include the ingress interface IP address, system IP address, and egress IP address of any of the hops being specified.

The **no** form of this command deletes hop list entries for the path. All LSPs currently using this path are affected. Additionally, all services actively using these LSPs are affected. The path must be shut down to delete the hop from the hop list. The **no hop hop-index** command will not result in any action except a warning message on the console indicating that the path is administratively up.

Parameters

hop-index

Specifies the hop index used to order the hops specified. The LSP always traverses from the lowest hop index to the highest. The hop index does not need to be sequential.

Values 1 to 1024

ip-address

Specifies the system or network interface IP address of the transit router. The IP address can be the interface IP address or the system IP address. If the system IP address is specified, the LSP can choose the best available interface. A hop list can also include the ingress interface IP address, system IP address, and egress IP address of any of the specified hops.

loose

Keyword to specify that the route taken by the LSP from the previous hop to this hop can traverse through other routers. Multiple hop entries with the same IP address are flagged as errors. Either the **loose** or **strict** keyword must be specified.

strict

Keyword to specify that the LSP must take a direct path from the previous hop router to this router. No transit routers between the previous router and this router are allowed. If the IP address specified is the interface address, that is the interface the LSP must use. If there are direct parallel links between the previous router and this router and if the system IP address is specified, any one of the available interfaces can be used by the LSP. The user must ensure that the previous router and this router have a direct link. Multiple hop entries with the same IP address are flagged as errors. Either the **loose** or **strict** keyword must be specified.

path

Syntax

[no] **path** *path-name*

Context

config>router>mpls

Platforms

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

Description

This command configures the path to be used for an LSP. A path can be used by multiple LSPs. A path can specify some or all hops from ingress to egress, and they can be either **strict** or **loose**. A path can also be empty (no path-name specified), in which case the LSP is set up based on the IGP (best effort) calculated shortest path to the egress router. Paths are created in a **shutdown** state. A path must be shut down before making any changes (adding or deleting hops) to the path. When a path is in the **shutdown** state, any LSP using the path becomes operationally down.

To create a strict path from the ingress to the egress router, the ingress and the egress routers must be included in the path statement.

The **no** form of this command deletes the path and all its associated configuration information. All the LSPs that are currently using this path will be affected. Additionally, all services that are actively using these LSPs will be affected. A path must be **shutdown** and unbound from all LSPs using the path before it can be deleted. The **no path path-name** command will not result in any action except a warning message on the console indicating that the path may be in use.

Parameters

path-name

Specifies a unique case-sensitive alphanumeric name label for the LSP path up to 32 characters in length.

shutdown

Syntax

[no] shutdown

Context

config>router>mpls>path

Platforms

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

Description

This command disables the existing LSPs using this path. All services using these LSPs are affected. Binding information, however, is retained in those LSPs. Paths are created in the **shutdown** state.

The **no** form of this command administratively enables the path. All LSPs, where this path is defined as primary or defined as standby secondary, are established or re-established.

Default

shutdown

2.17.2.1.9 Static LSP commands

static-lsp

Syntax

[no] static-lsp *lsp-name*

Context

config>router>mpls

Platforms

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

Description

This command configures a static LSP on the ingress router. The static LSP is a manually set up LSP where the next-hop IP address and the outgoing label (push) must be specified.

The **no** form of this command deletes this static LSP and associated information.

The LSP must be shut down to delete it. If the LSP is not shut down, the **no static-lsp *lsp-name*** command does nothing except generate a warning message on the console indicating that the LSP is administratively up.

Parameters

lsp-name

Specifies the name that identifies the LSP, up to 32 characters.

push

Syntax

push *label* **nexthop** *ip-address*

no push *label*

Context

config>router>mpls>static-lsp

Platforms

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

Description

This command configures the label to be pushed on the label stack and the next-hop IP address for the static LSP.

The **no** form of this command removes the association of the label to push for the static LSP.

Parameters

label

Specifies the label to push on the label stack. Label values 16 through 1,048,575 are defined as follows:

Label values 16 through 31 are reserved for the system.

Label values 32 through 1,023 are available for static assignment.

Label values 1,024 through 2,047 are reserved for future use.

Label values 2,048 through 18,431 are statically assigned for services.

Label values 28,672 through 131,071 are dynamically assigned for both MPLS and services.

Label values 131,072 through 1,048,575 are reserved for future use.

Values 16 to 1048575

nexthop *ip-address*

Specifies the IP address of the next hop toward the LSP egress router. If an ARP entry for the next hop exists, the static LSP is marked operational. If an ARP entry does not exist, the software sets the operational status of the static LSP to down and continues the ARP for the configured next hop. The Software continuously tries the ARP for the configured next hop at a fixed interval.

shutdown

Syntax

[no] shutdown

Context

config>router>mpls>static-lsp

Platforms

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

Description

This command administratively disables the static LSP.

The **no** form of this command administratively enables the static LSP.

Default

shutdown

to

Syntax

to *ip-address*

Context

config>router>mpls>static-lsp

Platforms

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

Description

This command specifies the system IP address of the egress router for the static LSP. When creating an LSP, this command is required. For LSPs that are used as transport tunnels for services, the **to** IP address must be the system IP address.

Parameters

ip-address

Specifies the system IP address of the egress router.

2.17.2.2 RSVP configuration commands

- [Generic commands](#)

- [RSVP commands](#)
- [Interface commands](#)
- [Message pacing commands](#)

2.17.2.2.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

config>router>rsvp

config>router>rsvp>interface

Platforms

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

Description

This command disables the RSVP protocol instance or the RSVP-related functions for the interface. The RSVP configuration information associated with this interface is retained. When RSVP is administratively disabled, all RSVP sessions are torn down. The existing configuration is retained.

The **no** form of this command administratively enables RSVP on the interface.

Default

shutdown

Special Cases

RSVP Protocol Handling on the 7210 SAS-Mxp

When the **no shutdown** command is issued in the **configure>router>rsvp** context, resources are allocated to enable the node to process the protocol. When the **configure>router>rsvp>shutdown** command is issued, the resources are deallocated.

2.17.2.2.2 RSVP commands

rsvp

Syntax

[no] rsvp

Context

config>router

Platforms

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

Description

Commands in this context configure RSVP protocol parameters. RSVP is not enabled by default and must be explicitly enabled (**no shutdown**).

RSVP is used to set up LSPs. RSVP should be enabled on all router interfaces that participate in signaled LSPs.

The **no** form of this command deletes this RSVP protocol instance and removes all configuration parameters for this RSVP instance. To suspend the execution and maintain the existing configuration, use the **shutdown** command. RSVP must be shut down before the RSVP instance can be deleted. If RSVP is not shut down, the **no rsvp** command does nothing except issue a warning message on the console indicating that RSVP is still administratively enabled.

Default

no shutdown

graceful-shutdown

Syntax

[no] graceful-shutdown

Context

config>router>rsvp

config>router>rsvp>interface

Platforms

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

Description

This command initiates a graceful shutdown of the specified RSVP interface or all RSVP interfaces on the node if applied at the RSVP level. These are referred to as maintenance interface and maintenance node, respectively.

To initiate a graceful shutdown the maintenance node generates a PathErr message with a specific error sub-code of Local Maintenance on TE Link required for each LSP that is exiting the maintenance interface.

The node performs a single make-before-break attempt for all adaptive CSPF LSPs it originates and LSP paths using the maintenance interfaces. If an alternative path for an affected LSP is not found, the LSP is maintained on its current path. The maintenance node also tears down and re-signals any detour LSP path using listed maintenance interfaces as soon as they are not active.

The maintenance node floods an IGP TE LSA/LSP containing Link TLV for the links under graceful shutdown with the traffic engineering metric set to 0xffffffff and unreserved bandwidth parameter set to zero (0).

After receiving the PathErr message, ahead-end LER node performs a single make-before-break attempt on the affected adaptive CSPF LSP. If an alternative path is not found, the LSP is maintained on its current path.

A node does not take any action on the paths of the following originating LSPs after receiving the PathErr message:

- an adaptive CSPF LSP for which the PathErr indicates a node address in the address list and the node corresponds to the destination of the LSP; in this case, no alternative paths can be found
- an adaptive CSPF LSP for which the path has explicit hops defined using the listed maintenance interfaces or nodes
- a CSPF LSP with the adaptive option disabled and for which the current path is over the listed maintenance interfaces in the PathErr message; these are not subject to make-before-break
- a non-CSPF LSP for which the current path is over the listed maintenance interfaces in the PathErr message

After receiving the updated IPG TE LSA/LSP for the maintenance interfaces, the head-end LER node updates the TE database. This information is used at the next scheduled CSPF computation for any LSP for which the path may traverse any of the maintenance interfaces.

The **no** form of this command disables the graceful shutdown operation at the RSVP interface level or at the RSVP level. The configured TE parameters of the maintenance links are restored and the maintenance node floods the links.

keep-multiplier

Syntax

keep-multiplier *number*

no keep-multiplier

Context

config>router>rsvp

Platforms

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

Description

This command configures the keep multiplier number. The **keep-multiplier** *number* is an integer used by RSVP to declare that a reservation is down or the neighbor is down.

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

Default

keep-multiplier 3

Parameters

number

Specifies the keep multiplier value.

Values 1 to 255

node-id-in-rro

Syntax

node-id-in-rro [**include** | **exclude**]

Context

config>router>rsvp

Platforms

7210 SAS-Mxp

Description

This command enables the option to include the node ID sub-object in the RRO. Propagation of the node ID sub-object is required to provide fast reroute protection for an LSP that spans multiple area domains.

If this option is disabled, the node ID is not included in the RRO object.

Default

node-id-in-rro exclude

Parameters

include

Keyword to include the node ID sub-object in the RRO.

exclude

Keyword to exclude the node ID sub-object in the RRO.

refresh-reduction-over-bypass

Syntax

refresh-reduction-over-bypass [**enable** | **disable**]

Context

config>router>rsvp

Platforms

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

Description

This command enables the refresh reduction capabilities over all bypass tunnels originating on this 7210 SAS PLR node or terminating on this 7210 SAS Merge Point (MP) node.

By default, this is disabled. Because a bypass tunnel may merge with the primary LSP path in a node downstream of the next hop, there is no direct interface between the PLR and the MP node, and it is possible the latter will not accept summary refresh messages received over the bypass.

When disabled, the node as a PLR or MP will not set the "Refresh-Reduction-Capable" bit on RSVP messages pertaining to LSP paths tunneled over the bypass. It will also not send the Message-ID in RSVP messages. This disables summary refresh.

Default

disable

Parameters

enable

Keyword to enable refresh reduction capabilities.

disable

Keyword to disable refresh reduction capabilities.

rapid-retransmit-time

Syntax

rapid-retransmit-time hundred-milliseconds

no rapid-retransmit-time

Context

config>router>rsvp

Platforms

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

Description

This command is used to define the value of the rapid retransmission interval. This is used in the retransmission mechanism based on the exponential back-off timer to handle unacknowledged message_id objects.

The RSVP message with the same message ID is retransmitted every $2 \times \text{rapid-retransmit-interval}$.

The node stops re transmission of unacknowledged RSVP messages whenever the updated back-off interval exceeds the value of the regular refresh interval or the number of retransmissions reaches the value of the rapid-retry-limit parameter, whichever comes first.

The rapid retransmission interval must be smaller than the regular refresh interval configured in **config>router>rsvp>refresh-time**.

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

Default

rapid-retransmit-time 5

Parameters

hundred-milliseconds

Specifies the rapid retransmission interval, in units of 100 milliseconds.

Values 1 to 100

rapid-retry-limit

Syntax

rapid-retry-limit *limit*

no rapid-retry-limit

Context

config>router>rsvp

Platforms

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

Description

This command configures the value of the rapid retry limit. This is used in the retransmission mechanism based on an exponential backoff timer to handle unacknowledged message_id objects. The RSVP message with the same message_id is retransmitted every 2 × rapid-retransmit-time interval of time. The node stops retransmission of unacknowledged RSVP messages whenever the updated backoff interval exceeds the value of the regular refresh interval or the number of retransmissions reaches the value of the **rapid-retry-limit** parameter, whichever comes first.

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

Default

rapid-retry-limit 3

Parameters

limit

Specifies the value of the rapid retry limit, expressed as an integer value.

Values 1 to 6

refresh-time

Syntax

refresh-time *seconds*

no refresh-time

Context

config>router>rsvp

Platforms

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

Description

This command controls the interval, in seconds, between the successive Path and Resv refresh messages. RSVP declares the session down after it misses **keep-multiplier** *number* consecutive refresh messages.

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

Default

refresh-time 30

Parameters

seconds

Specifies the refresh time in seconds.

Values 1 to 65535

2.17.2.2.3 Interface commands

interface

Syntax

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

Context

config>router>rsvp

Platforms

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

Description

This command enables RSVP protocol support on an IP interface. No RSVP commands are executed on an IP interface where RSVP is not enabled.

The **no** form of this command deletes all RSVP commands such as **hello-interval** and **subscription**, which are defined for the interface. The RSVP interface must be **shutdown** before it can be deleted. If the interface is not shut down, the **no interface ip-int-name** command does nothing except issue a warning message on the console indicating that the interface is administratively up.

Default

shutdown

Parameters

ip-int-name

Specifies the name of the network IP interface, up to 32 characters. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

authentication-key

Syntax

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

no authentication-key

Context

config>router>rsvp>interface

Platforms

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

Description

This command specifies the authentication key to be used between RSVP neighbors to authenticate RSVP messages. Authentication uses the MD-5 message-based digest.

When enabled on an RSVP interface, authentication of RSVP messages operates in both directions of the interface.

A node maintains a security association using one authentication key for each interface to a neighbor. The following items are stored in the context of this security association:

- HMAC-MD5 authentication algorithm
- key used with the authentication algorithm
- lifetime of the key (the user-entered key is valid until the user deletes it from the interface)
- source address of the sending system
- latest sending sequence number used with this key identifier

An RSVP sender transmits an authenticating digest of the RSVP message, computed using the shared authentication key and a keyed-hash algorithm. The message digest is included in an integrity object, which also contains a flags field, a key identifier field, and a sequence number field. The RSVP sender complies to the procedures for RSVP message generation as described in RFC 2747, *RSVP Cryptographic Authentication*.

An RSVP receiver uses the key together with the authentication algorithm to process received RSVP messages.

The MD5 implementation does not support the authentication challenge procedures as described in RFC 2747.

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

Default

no authentication-key

Parameters

authentication-key

Specifies the authentication key. The key can be any combination of ASCII characters up to 16 characters in length (unencrypted). If the string contains special characters (#, \$, 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 up to 33 alphanumeric characters. If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but for security purposes, the actual unencrypted key value is not provided.

hash

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.

bfd-enable

Syntax

[no] bfd-enable

Context

config>router>rsvp>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 RSVP interface. This causes RSVP to register the interface with the BFD session on that interface.

The user configures the BFD session parameters, such as **transmit-interval**, **receive-interval**, and **multiplier**, under the IP interface in the **config>router>interface>bfd** context.

The BFD session on the interface might already have been started because of a prior registration with another protocol; for example, OSPF or IS-IS.

The registration of an RSVP interface with BFD is performed when a neighbor gets its first session, which means registration occurs when this node sends or receives a new Path message over the interface. However, if the session did not come up because the session did not receive a RESV for a new Path message sent after the maximum number of retries, the LSP is shut down and the node deregisters with BFD. In general, the registration of RSVP with BFD is removed as soon as the last RSVP session is cleared.

The registration of an RSVP interface with BFD is performed independently of whether RSVP hello is enabled on the interface or not. However, hello timeout clears all sessions toward the neighbor and RSVP deregisters with BFD at the clearing of the last session.

An RSVP session is associated with a neighbor based on the interface address the Path message is sent to. If multiple interfaces exist to the same node, each interface is treated as a separate RSVP neighbor. The user must enable BFD on each interface, and RSVP will register with the BFD session running with each of those neighbors independently.

Similarly, disabling BFD on the interface results in removing registration of the interface with BFD.

When a BFD session transitions to the down state, the following actions are triggered. For RSVP signaled LSPs, this triggers activation of FRR bypass or detour backup LSPs (PLR role), global revertive (head-end role), and switchover to secondary (if any) (head-end role) for affected LSPs with FRR enabled. It triggers a switchover to secondary (if any) and scheduling of retries for signaling the primary path of the non-FRR affected LSPs (head-end role).

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

Default

no bfd-enable

hello-interval

Syntax

hello-interval *milli-seconds*

no hello-interval

Context

config>router>rsvp>interface

Platforms

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

Description

This command configures the time interval between RSVP hello messages.

RSVP hello packets are used to detect loss of RSVP connectivity with the neighboring node. Hello packets detect the loss of a neighbor more quickly than it would take for the RSVP session to time out based on the refresh interval. After the loss of the **keep-multiplier number** consecutive hello packets, the neighbor is declared to be in a down state.

The **no** form of this command reverts to the default value of the hello-interval. To disable sending hello messages, set the value to zero.

Default

hello-interval 3000

Parameters

milli-seconds

Specifies the RSVP hello interval in milliseconds, in multiples of 1000. A value of 0 (zero) disables the sending of RSVP hello messages.

Values 0 to 60000 milliseconds (in multiples of 1000)

implicit-null-label

Syntax

implicit-null-label [enable | disable]

no implicit-null-label

Context

config>router>rsvp

Platforms

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

Description

This command enables or disables the use of the implicit null label for all LSPs.

All LSPs for which this node is the egress LER and for which the path message is received from the previous hop node over this RSVP interface will signal the implicit null label. This means that if the egress LER is also the merge-point (MP) node, the incoming interface for the path refresh message over the bypass dictates if the packet uses the implicit null label or not. The same applies for a 1-to-1 detour LSP.

The user must shut down the RSVP interface before being able to change the implicit null configuration option.

The **no** form of this command resets the interface to the RSVP level configuration.

Default

implicit-null disable

Parameters

enable

Keyword to enable the implicit null label.

disable

Keyword to disable the implicit null label.

refresh-reduction

Syntax

[no] refresh-reduction

Context

config>router>rsvp>interface

Platforms

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

Description

This command enables the use of the RSVP overhead refresh reduction capabilities on this RSVP interface.

The 7210 SAS node accepts bundle RSVP messages from its peer over the interface, performs reliable RSVP message delivery to its peer, and uses summary refresh messages to refresh the path and resv states. Reliable message delivery must be explicitly enabled by the user after refresh reduction is enabled.

The other two capabilities are immediately enabled.

A bundle message reduces the overall message handling load; it consists of a bundle header followed by one or more bundle sub-messages. A bundle sub-message is any RSVP message other than a bundle message. A 7210 SAS node only processes the bundled RSVP messages received and does not generate them.

When reliable message delivery is supported by both the node and its peer over the RSVP interface, an RSVP message is sent with a message_id object. A message_id object can be added to any RSVP message, or it can be a sub-message of a bundled message.

If a node sets the ack_desired flag in the message_id object, the receiver acknowledges the receipt of the RSVP message by piggy-backing a message_ack object in the next RSVP message it sends to the node. Alternatively, an ACK message can also be used to send the message_ack object. In both cases, more than one message_ack object can be included in the same message.

The 7210 SAS supports only the use of ACK messages to send a message_ack object, but it can also process the received message_ack objects piggy-backed to hop-by-hop RSVP messages, such as Path and RESV.

The 7210 SAS sets the ack_desired flag only in non-refresh RSVP messages and in refresh messages that contain new state information.

A retransmission mechanism based on an exponential backoff timer is supported to handle unacknowledged message_id objects. An RSVP message with the same message_id is re-transmitted every 2× rapid-retransmit-time interval. The rapid-retransmit-time is referred to as the rapid

retransmission interval because it must be smaller than the regular refresh interval configured in the **config>router>rsvp>refresh-time** context.

The rapid retry limit indicates the maximum number of retransmissions allowed for unacknowledged RSVP messages. The node stops the retransmission of unacknowledged RSVP messages when:

- the updated backoff interval exceeds the regular refresh interval
- the number of retransmissions reaches the value of the rapid-retry-limit parameter, whichever comes first

These two parameters can be configured on a system in the **config>router>rsvp** context.

Summary refresh consists of sending a summary refresh messages containing message_id list objects. The fields of the message_id list object are populated with the values from the message_identifier field in the message_id object of a previously sent individual Path or RESV message. The summary refresh message is sent per refresh regular interval. The interval is configured by the user using the **refresh-time** command in the **config>router>rsvp** context. The receiver checks each message_id object against the saved Path and RESV states. If a match is found the state is updated. If any message_identifier field does not match, the node sends a message_id_nack object to the originator of the message.

The preceding capabilities are collectively referred to as "refresh overhead reduction extensions". When refresh-reduction is enabled on an RSVP interface, the node sets a "refresh-reduction-capable" bit in the flag field of the common RSVP header. If both peers on a RSVP interface set the "refresh-reduction-capable" bit, all the refresh overhead reduction extensions can be implemented. The node monitors the bit in all the RSVP messages received from the peer. The router stops sending summary refresh messages after the bit is cleared. the node does not send summary refresh messages if the bit is not set by the peer.

A node (with refresh reduction and reliable message delivery enabled) attempts to perform reliable message delivery even if the "refresh-reduction-capable" bit is not set by the peer. If the peer does not support the message_id object, it returns the "unknown object class" error message. The node retransmits the RSVP message without the message_id object and adopts the same message handling method for all future messages sent to the peer.

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

Default

no refresh-reduction

reliable-delivery

Syntax

[no] **reliable-delivery**

Context

config>router>rsvp>interface>refresh-reduction

Platforms

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

Description

This command enables reliable delivery of RSVP messages over the RSVP interface. When **refresh-reduction** is enabled on an interface and **reliable-delivery** is disabled, the router sends a message_id and not set ACK desired in the RSVP messages over the interface. Consequently, the 7210 SAS does not expect an ACK, but will accept it if received. The node also accepts message ID and reply with an ACK when requested. In this case, if the neighbor sets the "refresh-reduction-capable" bit in the flags field of the common RSVP header, the node enters summary refresh for a specific message_id it sent regardless of whether it received an ACK or not to this message from the neighbor.

When the **reliable-delivery** option is enabled on any interface, RSVP message pacing is disabled on all RSVP interfaces of the system, for example, the user cannot enable the **msg-pacing** option in the **config>router>rsvp** context, and an error message is returned in the CLI. Conversely, when the **msg-pacing** option is enabled, the user cannot enable the **reliable-delivery** option on any interface on this system. An error message is generated in the CLI after such an attempt.

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

Default

no reliable-delivery

subscription

Syntax

subscription *percentage*

no subscription

Context

config>router>rsvp>interface

Platforms

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

Description

This command configures the percentage of the link bandwidth that RSVP can use for reservation and sets a limit for the amount of over-subscription or under-subscription allowed on the interface.

When the **subscription** command is set to zero, no new sessions are permitted on this interface. If the *percentage* value is exceeded, the reservation is rejected and a log message is generated.

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

Default

subscription 100

Parameters

percentage

Specifies the percentage of the interface bandwidth that RSVP allows to be used for reservations.

Values 0 to 1000

2.17.2.2.4 Message pacing commands

msg-pacing

Syntax

[no] msg-pacing

Context

config>router>rsvp

Platforms

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

Description

This command enables RSVP message pacing for which the specified number of RSVP messages, specified in the **max-burst** command, are sent in a configured interval, specified in the **period** command. A count is kept of the messages that were dropped because the output queue for the interface used for message pacing was full.

Default

no msg-pacing

max-burst

Syntax

max-burst *number*

no max-burst

Context

config>router>rsvp>msg-pacing

Platforms

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

Description

This command configures the maximum number of RSVP messages that are sent in the specified period under normal operating conditions.

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

Default

max-burst 650

Parameters

number

Specifies the maximum number of RSVP messages sent within the specified period.

Values 100 to 1000 (in increments of 10)

period

Syntax

period *milli-seconds*

no period

Context

config>router>rsvp>msg-pacing

Platforms

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

Description

This command configures the time interval, in milliseconds, during which the router can send RSVP messages, as specified in the **max-burst** command.

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

Default

period 100

Parameters

milli-seconds

Specifies the time interval for sending RSVP messages.

Values 100 to 1000 milliseconds (in increments of 10 milliseconds)

2.17.2.3 Show commands

- [Show MPLS commands](#)
- [Show RSVP commands](#)

2.17.2.3.1 Show MPLS commands

admin-group

Syntax
`admin-group group-name`

Context
show>router>if-attribute

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

Description
This command displays MPLS administrative group information.

Parameters
group-name
Specifies a group name, up to 32 characters.

Output
The following output is an example of MPLS administrative group information, and [Table 10: Output fields: If-attribute admin group](#) describes the output fields.

Sample output

```
A:ALA-1# show router if-attribute admin-group
=====
MPLS Administrative Groups
=====
Group Name                Group Value
-----
green                      15
red                        25
yellow                     20
-----
No. of Groups: 3
=====
A:ALA-1#
```

Table 10: Output fields: If-attribute admin group

Label	Description
Group Name	Displays the name of the group; the name identifies the administrative group within a virtual router instance

Label	Description
Group Value	Displays the unique group value associated with the administrative group. If the value displays -1, then the group value for this entry has not been set.
No. of Groups	Displays the total number of configured admin groups within the virtual router instance

bypass-tunnel

Syntax

bypass-tunnel [*to ip-address*] [**protected-lsp** [*lsp-name*]] [**dynamic** | **manual** | **p2mp**] [**detail**]

Context

show>router>mpls

Platforms

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

Description

This command displays information about bypass tunnels.

If fast reroute is enabled on an LSP and the facility method is selected, instead of creating a separate LSP for every LSP that is to be backed up, a single LSP is created which serves as a backup for a set of LSPs. Such an LSP tunnel is called a bypass tunnel.

Parameters

ip-address

Specifies the IP address of the egress router.

lsp-name

Specifies the name of the LSP protected by the bypass tunnel.

dynamic

Keyword to display dynamically assigned labels for bypass protection.

manual

Keyword to display manually assigned labels for bypass protection.

p2mp

Keyword to display P2MP bypass tunnel information.

detail

Keyword to display detailed information.

Output

The following output is an example of MPLS bypass tunnel information, and [Table 11: Output fields: MPLS bypass-tunnel](#) describes the output fields.

Sample output

```
*A:Dut-A>config>router>mpls# show>router>mpls# bypass-tunnel

*A:SRU4>show>router>mpls# bypass-tunnel
=====
MPLS Bypass Tunnels
=====
Legend : m - Manual d - Dynamic p - P2mp
=====
To State Out I/F Out Label Reserved Protected Type
BW (Kbps) LSP Count
-----
No Matching Entries Found
=====

*A:SRU4>show>router>mpls#*A:Dut-A>show>router>mpls#

*A:Dut-A>config>router>mpls# show>router>mpls# bypass-tunnel detail
=====
MPLS Bypass Tunnels (Detail)
=====
-----
bypass-node10.20.1.2
-----
To           : 10.20.1.4           State          : Up
Out I/F      : 1/1/2              Out Label      : 131070
Up Time     : 0d 00:00:18         Active Time    : n/a
Reserved BW  : 0 Kbps             Protected LSP Count : 1
Type        : Dynamic
Setup Priority : 7                 Hold Priority   : 0
Class Type   : 0
Exclude Node : None               Inter-Area      : False
ComputedHops:
  10.20.1.1, If Index : 2(S)
  -> 10.20.1.2, If Index : 2(S)
  -> 10.20.1.4, If Index : 2(S)
  -> 10.20.1.6, If Index : 2(S)
Actual Hops :
  10.20.1.1, If Index: 2 @ n           Record Label : N/A
  -> 10.20.1.2, If Index : 2 @ n       Record Label : 131071
  -> 10.20.1.4, If Index : 2           Record Label : 131071
  -> 10.20.1.6, If Index : 2           Record Label : 131071
=====
*A:Dut-A>show>router>mpls#
```

Table 11: Output fields: MPLS bypass-tunnel

Label	Description
To	Displays the system IP address of the egress router
State	Dispalys the LSP administrative state

Label	Description
Out I/F	Displays the name of the network IP interface
Out Label	Displays the incoming MPLS label on which to match
Reserved BW (Kbps)	Displays the amount of bandwidth in megabits per second (Mbps) reserved for the LSP

interface

Syntax

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

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

Context

show>router>mpls

Platforms

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

Description

This command displays MPLS interface information.

Parameters

ip-int-name

Specifies the name of the network IP interface. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

ip-address

Displays the system or network interface IP address.

label-map label

Displays the MPLS label on which to match.

Values 32 to 1048575

Output

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

Sample output

```
A:7210SAS# config>router>mpls# show router mpls interface
```

```
=====
MPLS Interfaces
```



```

=====
Interface                Port-id          Adm   Opr   TE-metric
-----
system                   system           Up    Up    None
  Admin Groups           None
  Srlg Groups            None
ip-10.10.2.3             1/1/15          Up    Up    None
  Admin Groups           None
  Srlg Groups            None
ip-10.10.5.3             1/1/1           Up    Up    None
  Admin Groups           None
  Srlg Groups            None
ip-10.10.11.3            1/1/3           Up    Up    None
  Admin Groups           None
  Srlg Groups            None
ip-10.10.12.3            lag-1           Up    Up    None
  Admin Groups           None
  Srlg Groups            None
-----
Interfaces : 5
=====
*A:7210SAS#

```

Table 12: Output fields: MPLS interface

Label	Description
Interface	Displays the interface name
Port-id	Displays the port ID displayed in the <i>slot/mda/port</i> format
Adm	Displays the administrative state of the interface
Opr	Displays the operational state of the interface
Srlg Groups	Displays the shared risk link group (SRLG) names
Te-metric	Displays the traffic engineering metric used on the interface
Interfaces	Displays the total number of interfaces
Transmitted	Displays the number of packets and octets transmitted from the interface
Received	Displays the number of packets and octets received
In Label	Displays the ingress label
In I/F	Displays the ingress interface
Out Label	Displays the egress label
Out I/F	Displays the egress interface
Next Hop	Displays the next hop IP address for the static LSP
Type	Displays whether the label value is statically or dynamically assigned

label

Syntax

label start-label [end-label | in-use | label-owner]

Context

show>router>mpls-labels

Platforms

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

Description

This command displays information about MPLS labels exchanged by signaling protocols.

Parameters

start-label

Specifies the label value assigned at the ingress router.

Values 32 to 131071

end-label

Specifies the label value assigned for the egress router.

Values 32 to 131071

in-use

Specifies the number of in-use labels displayed.

label-owner

Specifies the owner of the label.

Values bgp | ildp | mirror | rsvp | static | sr | svcmgr | tldp | vprn

Output

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

Sample output

```
*A:SRU4>config>router>mpls-labels# show router mpls-labels label 202
=====
MPLS Label 202
=====
Label          Label Type      Label Owner
-----
202            static-lsp      STATIC
-----
In-use labels in entire range          : 5057
=====
```

```
*A:SRU4>config>router>mpls-labels#
```

Table 13: Output fields: MPLS label

Label	Description
Label	Displays the label value
Label Type	Displays whether the label value is statically or dynamically assigned
Label Owner	Displays the label owner
In-use labels in entire range	Displays the total number of labels being used by RSVP

label-range

Syntax

label-range

Context

show>router>mpls-labels

Platforms

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

Description

This command displays the MPLS label range.

Output

The following output is an example of MPLS label range information, and [Table 14: Output fields: MPLS label-range](#) describes the output fields.

Sample output

```
*A:Dut-A# show router mpls-labels label-range
=====
Label Ranges
=====
Label Type      Start Label      End Label      Aging      Total Available
-----
static-lsp      32               1023           -          992
static-svc      2048             18431          -          16384
dynamic         32768            131071         0          102400
=====
*A:Dut-A#
```

Table 14: Output fields: MPLS label-range

Label	Description
Label Type	Displays information about the static-lsp , static-svc , and dynamic label types.
Start Label	Displays the label value assigned at the ingress router
End Label	Displays the label value assigned for the egress router
Aging	Displays the number of labels released from a service that are transitioning back to the label pool; labels are aged 15 seconds
Total Available	Displays the number of label values available

lsp

Syntax

lsp *lsp-name* [status {up|down}] [from *ip-address* | to *ip-address*] [detail]

lsp {transit | terminate} [status {up | down}] [from *ip-address* | to *ip-address* | **lsp-name** *name*] [detail]

lsp count

lsp *lsp-name* **activepath**

lsp *lsp-name* **path** [*path-name*] [status {up | down}] [detail]

lsp [*lsp-name*] **path** [*path-name*] **mbb**

Context

show>router>mpls

Platforms

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

Description

This command displays LSP details.

Parameters

lsp lsp-name

The name of the LSP used in the path.

status up

Keyword to display an LSP that is operationally up.

status down

Keyword to display an LSP that is operationally down.

- from ip-address**
Displays the IP address of the ingress router for the LSP.
- to ip-address**
Displays the IP address of the egress router for the LSP.
- transit**
Displays the number of static LSPs that transit through the router.
- terminate**
Displays the number of static LSPs that terminate at the router.
- lsp count**
Displays the total number of LSPs.
- activepath**
Displays the present path being used to forward traffic.
- mbb**
Displays make-before-break (MBB) information.
- detail**
Displays detailed information.

Output

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

Sample output

```
*A:SRU4>config>router>mpls#  show router mpls lsp "to_110_20_1_1_cspf"
=====
MPLS LSPs (Originating)
=====
LSP Name                               To                               Fastfail   Adm   Opr
                                Config
-----
to_110_20_1_1_cspf                   110.20.1.1                       No         Up    Up
-----
LSPs : 1
=====

*A:SRU4>config>router>mpls#

*A:Dut-A# show router mpls lsp transit detail
=====
MPLS LSPs (Transit) (Detail)
=====
LSP D_B_1::D_B_1
-----
From           : 10.20.1.4           To           : 10.20.1.2
State          : Up
In Interface   : lag-1:10           In Label     : 130668
Out Interface  : lag-2              Out Label    : 131065
Previous Hop   : 10.10.14.4         Next Hop     : 10.10.12.2
Reserved BW    : 0 Kbps
-----
```

```
*A:Dut-A#

*=====
*A:7210-SAS>show>router>mpls# lsp A detail

=====
MPLS LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name      : A                      LSP Tunnel ID : 1
From          : 10.2.2.2              To            : 10.100.100.100
Adm State     : Up                    Oper State     : Down
LSP Up Time   : 0d 00:00:00           LSP Down Time  : 0d 00:05:42
Transitions   : 2                     Path Changes   : 2
Retry Limit   : 0                     Retry Timer    : 30 sec
Signaling     : RSVP                  Resv. Style    : SE
Hop Limit     : 255                   Negotiated MTU : 0
Adaptive      : Enabled                ClassType      : 0
FastReroute   : Disabled              Oper FR        : Disabled
CSPF          : Disabled               ADSPEC         : Disabled
Metric        : 0
Include Grps:                          Exclude Grps  :
None                                                  None
Type          : RegularLsp             Least Fill    : Disabled
LdpOverRsvp   : Enabled                VprnAutoBind  : Enabled
Oper Metric   : 65535

Primary       : A                      Down Time     : 0d 00:05:42
Bandwidth     : 0 Mbps
=====
*A:7210-SAS>show>router>mpls# lsp 2 detail

*A:Dut-A# config>router>mpls# show router mpls lsp "1" path detail
=====
MPLS LSP 1 Path (Detail)
=====
Legend :
  @ - Detour Available      # - Detour In Use
  b - Bandwidth Protected  n - Node Protected
  s - Soft Preemption
  S - Strict                L - Loose
  A - ABR
=====
-----
LSP 1 Path 1
-----
LSP Name      : 1                      Path LSP ID   : 30208
From          : 10.20.1.1              To            : 10.20.1.6
Adm State     : Up                    Oper State     : Up
Path Name     : 1                     Path Type      : Primary
Path Admin    : Up                    Path Oper      : Up
OutInterface  : 1/1/1                 Out Label     : 131071
Path Up Time  : 0d 00:00:05           Path Dn Time   : 0d 00:00:00
Retry Limit   : 0                     Retry Timer    : 30 sec
RetryAttempt  : 0                     NextRetryIn    : 0 sec

Adspec        : Disabled               Oper Adspec    : Disabled
CSPF          : Enabled                Oper CSPF      : Enabled
Least Fill    : Disabled               Oper LeastF*   : Disabled
```

```

FRR      : Enabled
FRR NodePro*: Enabled
FR Hop Limit: 16
FR Prop Adm*: Disabled
Prop Adm Grp: Disabled
Inter-area : False

Neg MTU    : 1496
Bandwidth  : No Reservation
Hop Limit  : 255
Record Route: Record
Record Label: Record
SetupPriori*: 7
Hold Priori*: 0
Class Type : 0
Backup CT  : None
MainCT Retry: n/a
  Rem      :
MainCT Retry: 0
  Limit    :
Include Grps:
None
Exclude Grps:
None

Adaptive    : Enabled
Preference  : n/a
Path Trans  : 1
Failure Code: noError
ExplicitHops:
  No Hops Specified
Actual Hops :
  10.20.1.1, If Index : 2 @ n
N/A
-> 10.20.1.2, If Index : 2 @ n
131071
-> 10.20.1.4, If Index : 2
131071
-> 10.20.1.6, If Index : 2
131071
ComputedHops:
  10.20.1.1, If Index : 2(S)
  -> 10.20.1.2, If Index : 2(S)
  -> 10.20.1.4, If Index : 2(S)
  -> 10.20.1.6, If Index : 2(S)
ResigEligib*: False
LastResignal: n/a

Oper FRR      : Enabled
Oper FRR NP   : Enabled
Oper FRHopL*: 16
Oper FRProp*: Disabled
Oper PropAG   : Disabled

Oper MTU      : 1496
Oper Bw       : 0 Mbps
Oper HopLim*: 255
Oper RecRou*: Record
Oper RecLab*: Record
Oper SetupP*: 7
Oper HoldPr*: 0
Oper CT       : 0

Oper InclGr*:
None
Oper ExclGr*:
None

Oper Metric : 3000

CSPF Queries: 1
Failure Node: n/a

Record Label      :
Record Label      :
Record Label      :
Record Label      :

CSPF Metric : 3000
=====
* indicates that the corresponding row element may have been truncated.

```

Table 15: Output fields: MPLS LSP

Label	Description
LSP Name	Displays the name of the LSP used in the path
To	Displays the system IP address of the egress router for the LSP
Adm State	Down — the path is administratively disabled Up — the path is administratively enabled

Label	Description
Oper State	Down — the path is operationally down Up — the path is operationally up
Oper State	Down — the path is operationally down Up — the path is operationally up
LSPs	Displays the total number of LSPs configured
From	Displays the IP address of the ingress router for the LSP
LSP Up Time	Displays the length of time the LSP has been operational
Transitions	Displays the number of transitions that have occurred for the LSP
Retry Limit	Displays the number of attempts that the software should make to re-establish the LSP after it has failed
Signaling	Displays the signaling style
Hop Limit	Displays the maximum number of hops that an LSP can traverse, including the ingress and egress routers
Fast Reroute/FastFail Config	enabled — Fast reroute is enabled. In the event of a failure, traffic is immediately rerouted on the precomputed detour LSP, thus minimizing packet loss. disabled — there is no detour LSP from each node on the primary path
ADSPEC	enabled — the LSP will include advertising data (ADSPEC) objects in RSVP messages disabled — the LSP will not include advertising data (ADSPEC) objects in RSVP messages
Primary	Displays the preferred path for the LSP
Secondary	Displays the alternate path that the LSP uses if the primary path is not available.
Bandwidth	Displays the amount of bandwidth in megabits per second (Mbps) reserved for the LSP path
LSP Up Time	Displays the total time in increments that the LSP path has been operational
LSP Tunnel ID	Displays the value that identifies the label switched path that is signaled for this entry
To	Displays the IP address of the egress router for the LSP

Label	Description
LSP Down Time	Displays the total time in increments that the LSP path has not been operational
Path Changes	Displays the number of path changes this LSP has had. For every path change (path down, path up, path change), a corresponding syslog/trap (if enabled) is generated.
Retry Timer	Displays the time, in seconds, for LSP re-establishment attempts after an LSP failure
Resv Style	se — Specifies a shared reservation environment with a limited reservation scope. This reservation style creates a single reservation over a link that is shared by an explicit list of senders. ff — Specifies a shared reservation environment with an explicit reservation scope. Specifies an explicit list of senders and a distinct reservation for each of them.
Negotiated MTU	Displays the size of the maximum transmission unit (MTU) that is negotiated during establishment of the LSP
FR Hop Limit	Displays the total number of hops a detour LSP can take before merging back onto the main LSP path
LastResignalAttempt	Displays the system up time when the last attempt to resignal this LSP was made
VprnAutoBind	Displays the status on VPRN auto-bind feature as enabled or disabled

oam-template

Syntax

oam-template [*template-name*] [**associations**]

Context

show>router>mpls>mpls-tp

Platforms

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

Description

This command displays MPLS-TP OAM template information.

Output

The following output is an example of MPLS-TP OAM template information.

Sample output

```
*A:7210SAS>show>router>mpls>tp# oam-template

=====
MPLS-TP OAM Templates
=====
Template Name : temp1          Router ID      : 1
BFD Template  : temp1          Hold-Down Time: 0 centiseconds
                                   Hold-Up Time   : 0 deciseconds
-----
Template Name : temp2          Router ID      : 1
BFD Template  : temp2          Hold-Down Time: 0 centiseconds
                                   Hold-Up Time   : 0 deciseconds
-----
Template Name : temp3          Router ID      : 1
BFD Template  : temp3          Hold-Down Time: 0 centiseconds
                                   Hold-Up Time   : 0 deciseconds
=====
*A:7210SAS>show>router>mpls>tp#
```

protection-template

Syntax

protection-template [*template-name*] [*associations*]

Context

show>router>mpls>mpls-tp

Platforms

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

Description

This command displays MPLS-TP protection template information.

Output

The following output is an example of MPLS-TP protection template information.

Sample output

```
*A:7210SAS>show>router>mpls>tp# protection-template

=====
MPLS-TP Protection Templates
=====
Template Name : temp1          Router ID      : 1
Protection Mode: one2one       Direction      : bidirectional
Revertive      : revertive     Wait-to-Restore : 1sec
Rapid-PSC-Timer: 10ms         Slow-PSC-Timer  : 5sec
-----
Template Name : temp2          Router ID      : 1
Protection Mode: one2one       Direction      : bidirectional
```

```

Revertive      : revertive          Wait-to-Restore : 1sec
Rapid-PSC-Timer: 10ms              Slow-PSC-Timer  : 5sec
-----
Template Name  : temp3              Router ID       : 1
Protection Mode: one2one            Direction       : bidirectional
Revertive      : revertive          Wait-to-Restore : 1sec
Rapid-PSC-Timer: 10ms              Slow-PSC-Timer  : 5sec
-----
=====
*A:7210SAS>show>router>mpls>tp#

```

status

Syntax

status

Context

show>router>mpls>mpls-tp

Platforms

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

Description

This command displays MPLS-TP system configuration information.

Output

The following output is an example of MPLS-TP system configuration information.

Sample output

```

*A:mlstp-dutA# show router mpls mpls-tp status
=====
MPLS-TP Status
=====
Admin Status   : Up
Global ID      : 42
Tunnel Id Min  : 1
Node ID        : 0.0.3.233
Tunnel Id Max  : 4096
=====

```

transit-path

Syntax

transit-path [*path-name*] [*detail*]

Context

show>router>mpls>mpls-tp

Platforms

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

Description

This command displays MPLS-TP tunnel information.

Parameters

- path-name**
Specifies the path name, up to 32 characters max.
- detail**
Keyword to displays detailed information.

Output

The following output is an example of MPLS-TP tunnel information.

Sample output

```
A:mplstp-dutC# show router mpls mpls-tp transit-path
<path-name>
"tp-32"  "tp-33"  "tp-34"  "tp-35"  "tp-36"  "tp-37"  "tp-38"  "tp-39"
"tp-40"  "tp-41"
detail

A:mplstp-dutC# show router mpls mpls-tp transit-path "tp-32"

=====
MPLS-TP Transit tp-32 Path Information
=====
Path Name      : tp-32
Admin State    : Up
Oper State     : Up

-----
Path      NextHop      InLabel  OutLabel  Out I/F
-----
FP                2080      2081      CtoB_1
RP                2081      2080      CtoA_1
=====

A:mplstp-dutC# show router mpls mpls-tp transit-path "tp-32" detail

=====
MPLS-TP Transit tp-32 Path Information (Detail)
=====
Path Name      : tp-32
Admin State    : Up
Oper State     : Up

-----
Path ID configuration
Src Global ID  : 42
Src Node ID   : 0.0.3.234
LSP Number    : 2
Dst Global ID  : 42
Dst Node ID   : 0.0.3.233
Dst Tunnel Num: 32

Forward Path configuration
In Label       : 2080
Out Interface  : CtoB_1
Out Label      : 2081
Next Hop Addr  : n/a

Reverse Path configuration
In Label       : 2081
Out Interface  : CtoA_1
Out Label      : 2080
Next Hop Addr  : n/a
```

```
=====
A:mplstp-dutC#
```

srlg-database

Syntax

srlg-database [**router-id** *ip-address*] [**interface** *ip-address*]

Context

show>router>mpls

Platforms

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

Description

This command displays MPLS SRLG database information.

Parameters

router-id *ip-address*

Specifies a 32-bit integer uniquely identifying the router in the autonomous system. By convention, to ensure uniqueness, this may default to the value of one of the router IPv4 host addresses, represented as a 32-bit unsigned integer, if IPv4 is configured on the router. The router-id can be either the local one or a remote router.

interface *ip-address*

Specifies the IP address of the interface.

path

Syntax

path [*path-name*] [**lsp-binding**]

Context

show>router>mpls

Platforms

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

Description

This command displays MPLS paths.

Parameters

path-name

Specifies the unique name label for the LSP path, up to 32 characters.

lsp-binding

Keyword to display binding information.

Output

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

Sample output

```
*A:SRU4>config>router>mpls# show router mpls path
=====
MPLS Path:
=====
Path Name                Adm  Hop Index  IP Address  Strict/Loose
-----
to_110_20_1_1            Up   no hops    n/a         n/a
to_110_20_1_2            Up   no hops    n/a         n/a
to_110_20_1_3            Up   no hops    n/a         n/a
to_110_20_1_4            Up   no hops    n/a         n/a
to_110_20_1_5            Up   no hops    n/a         n/a
to_110_20_1_6            Up   no hops    n/a         n/a
to_110_20_1_110          Up   no hops    n/a         n/a
to_10_8_100_15           Up   no hops    n/a         n/a
to_10_20_1_20            Up   no hops    n/a         n/a
to_10_20_1_22            Up   no hops    n/a         n/a
to_10_100_1_1            Up   no hops    n/a         n/a
-----
Paths : 11
=====
*A:SRU4>config>router>mpls#
```

Table 16: Output fields: MPLS path

Label	Description
Path Name	Displays the unique name label for the LSP path
Adm	Down — the path is administratively disabled Up — the path is administratively enabled
Hop Index	Displays the value used to order the hops in a path

Label	Description
IP Address	Displays the IP address of the hop that the LSP should traverse on the way to the egress router
Strict/Loose	Strict — the LSP must take a direct path from the previous hop router to the next router Loose — the route taken by the LSP from the previous hop to the next hop can traverse through other routers
LSP Name	Displays the name of the LSP used in the path
Binding	Primary — the preferred path for the LSP Secondary — the standby path for the LSP
Paths	Displays the total number of paths configured

srlg-group

Syntax

srlg-group [*group-name*]

Context

show>router>mpls

Platforms

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

Description

This command displays MPLS SRLG group information.

Parameters

group-name

Specifies the name of the SRLG group within a virtual router instance.

Output

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

Sample output

```
*A:SRU4>config>router>mpls# show router mpls srlg-group
=====
MPLS Srlg Groups
=====
Group Name                Group Value    Interfaces
-----
1432                      1432          srl-1
```

```

1433          1433          srl-3
1434          1434          aps-8
1435          1435          aps-9
2410          2410          srr-1
2411          2411          srr-2
2412          2412          srr-3
3410          3410          aps-1
3420          3420          aps-2
3430          3430          aps-3
3440          3440          sr4-1
41.80         4180         g7600
41104         41104        germ-1
415.70        41570        gsr1
420.40        42040        m160
422.60        42260        gsr2
44.200        44200        hubA
45100         45100        ess-7-1
45110         45110        ess-7-2
45120         45120        ess-7-3
4651          4651          src-1.1
4652          4652          src-1.2
4653          4653          src-1.3
4654          4654          src-1.4
4655          4655          src-1.5
4656          4656          src-1.6
4657          4657          src-1.7
4658          4658          src-1.8
4659          4659          src-1.9
4660          4660          src-1.10

```

```
-----
No. of Groups: 30
=====
```

```
*A:SRU4>config>router>mpls#
```

```
*A:SRU4>config>router>mpls# show router mpls srlg-group "1432"
```

```
=====
MPLS Srlg Groups
=====
```

```

Group Name          Group Value    Interfaces
-----
1432                1432          srl-1
-----

```

```
No. of Groups: 1
=====
```

```
*A:SRU4>config>router>mpls#
```

Table 17: Output fields: MPLS SRLG-group

Label	Description
Group Name	Displays the name of the SRLG group within a virtual router instance
Group Value	Displays the group value associated with this SRLG group
Interface	Displays the interface where the SRLG group is associated
No. of Groups	Displays the total number of SRLG groups associated with the output

static-lsp

Syntax

static-lsp [/sp-name]
static-lsp {transit | terminate}
static-lsp count

Context

show>router>mpls

Platforms

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

Description

This command displays MPLS static LSP information.

Output

The following output is an example of MPLS static LSP information, and [Table 18: Output fields: MPLS static-LSP](#) describes the output fields.

Sample output

```
A:ALA-12# show router mpls static-lsp
=====
MPLS Static LSPs (Originating)
=====
Lsp Name          To          Next Hop      Out Label Out I/F  Adm  Opr
-----
NYC_SJC_customer2 10.20.1.10  10.10.1.4     1020      1/1/1    Up   Up
-----
LSPs : 1
=====

A:ALA-12#

*A:SRU4>config>router>mpls# show router mpls static-lsp transit
=====
MPLS Static LSPs (Transit)
=====
In Label   In Port   Out Label   Out Port   Next Hop      Adm  Opr
-----
240        aps-1     440         1/1/10     10.22.11.3    Up   Up
241        aps-1     441         1/1/10     10.22.11.3    Up   Up
242        aps-1     442         1/1/10     10.22.11.3    Up   Up
243        aps-1     443         1/1/10     10.22.11.3    Up   Up
244        aps-1     444         1/1/10     10.22.11.3    Up   Up
245        aps-1     445         1/1/10     10.22.11.3    Up   Up
246        aps-1     446         1/1/10     10.22.11.3    Up   Up
247        aps-1     447         1/1/10     10.22.11.3    Up   Up
248        aps-1     448         1/1/10     10.22.11.3    Up   Up
249        aps-1     449         1/1/10     10.22.11.3    Up   Up
250        aps-1     450         1/1/10     10.22.11.3    Up   Up
251        aps-1     451         1/1/10     10.22.11.3    Up   Up
```

```

252      aps-1      452      1/1/10      10.22.11.3      Up      Up
253      aps-1      453      1/1/10      10.22.11.3      Up      Up
...
207      3/2/8      407      1/1/9      10.22.10.3      Up      Up
208      3/2/8      408      1/1/9      10.22.10.3      Up      Up
209      3/2/8      409      1/1/9      10.22.10.3      Up      Up
-----
LSPs : 256
=====
*A:SRU4>config>router>mpls#

A:ALA-12# show router mpls static-lsp terminate
=====
MPLS Static LSPs (Terminate)
=====
In Label    In I/F      Out Label   Out I/F     Next Hop    Adm   Opr
-----
1021        1/1/1      n/a         n/a         n/a         Up    Up
-----
LSPs : 1
=====
A:ALA-12#

```

Table 18: Output fields: MPLS static-LSP

Label	Description
Lsp Name	Displays the name of the LSP used in the path
To	Displays the system IP address of the egress router for the LSP
Next Hop	Displays the system IP address of the next hop in the LSP path
In I/F	Displays the ingress interface
Out Label	Displays the egress label
Out I/F	Displays the egress interface
Adm	Down — the path is administratively disabled Up — the path is administratively enabled
Opr	Down — the path is operationally down Up — the path is operationally up
LSPs	Displays the total number of static LSPs

status

Syntax

status

Context

show>router>mpls

Platforms

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

Description

This command displays MPLS operation information.

Output

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

Sample output

```

*A:7210SAS# show router mpls status

=====
MPLS Status
=====
Admin Status      : Up                Oper Status      : Up
Oper Down Reason  : n/a
FR Object         : Enabled           Resignal Timer   : Disabled
Hold Timer        : 1 seconds         Next Resignal    : N/A
Srlg Frr          : Disabled          Srlg Frr Strict  : Disabled
Dynamic Bypass    : Enabled           User Srlg Database : Disabled
Least Fill Min Thd.: 5 percent        LeastFill ReoptiThd: 10 percent
Short. TTL Prop Lo*: Enabled          Short. TTL Prop Tr*: Enabled
AB Sample Multipli*: 1                AB Adjust Multipli*: 288
Exp Backoff Retry : Disabled          CSPF On Loose Hop : Disabled
Lsp Init RetryTime*: 30 seconds
Logger Event Bundl*: Disabled

P2mp Resignal Timer: Disabled          P2mp Next Resignal : N/A
Sec FastRetryTimer : Disabled          Static LSP FR Timer: 30 seconds
P2P Max Bypass Ass*: 1000
P2PActPathFastRetry: Disabled
In Maintenance Mode: No

LSP Counts      Originate      Transit      Terminate
-----
Static LSPs     0                0            0
Dynamic LSPs    0                0            1
Detour LSPs     0                0            0
S2l LSPs        0                0            0
=====
* indicates that the corresponding row element may have been truncated.
*A:7210SAS#

```

Table 19: Output fields: MPLS status

Label	Description
Admin Status	Down — MPLS is administratively disabled Up — MPLS is administratively enabled

Label	Description
Oper Status	Down — MPLS is operationally down Up — MPLS is operationally up
LSP Counts	Static LSPs — Displays the count of static LSPs that originate, transit, and terminate on or through the router Dynamic LSPs — Displays the count of dynamic LSPs that originate, transit, and terminate on or through the router Detour LSPs — Displays the count of detour LSPs that originate, transit, and terminate on or through the router
FR Object	Enabled — Specifies that Fast reroute object is signaled for the LSP Disabled — Specifies that Fast reroute object is not signaled for the LSP
Resignal Timer	Enabled — Specifies that the resignal timer is enabled for the LSP Disabled — Specifies that the resignal timer is disabled for the LSP
Hold Timer	Displays the amount of time that the ingress node holds before programming its data plane and declaring the LSP up to the service module

2.17.2.3.2 Show RSVP commands

interface

Syntax

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

Context

show>router>rsvp

Platforms

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

Description

This command shows RSVP interfaces.

Parameters

ip-int-name

Specifies the name of the network IP interface, up to 32 characters. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

ip-address

Specifies the system or network interface IP address.

statistics

Keyword to display the IP address and number of packets sent and received on an interface-basis.

detail

Keyword to display detailed information.

Output

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

Sample output

```
*A:7210-SAS>show>router>rsvp# interface detail

=====
RSVP Interfaces (Detailed)
=====
-----
Interface : system
-----
Interface      : system
Port ID       : system
Admin State    : Up
Active Sessions : 0
Total Sessions : 0
Subscription   : 100 %
Total BW       : 0 Mbps
Hello Interval : n/a
Authentication : Disabled
Auth Rx Seq Num : n/a
Auth Tx Seq Num : n/a
Refresh Reduc. : Disabled
Bfd Enabled    : n/a
ImplicitNullLabel : Disabled*
Oper State     : Up
Active Resvs   : 0
Port Speed     : 0 Mbps
Aggregate      : Dsabl
Hello Timeouts : n/a
Auth Key Id    : n/a
Auth Win Size  : n/a
Reliable Deli. : Disabled
Graceful Shut. : Disabled
GR helper      : n/a

Percent Link Bandwidth for Class Types
Link Bw CT0    : 100
Link Bw CT1    : 0
Link Bw CT2    : 0
Link Bw CT3    : 0
Link Bw CT4    : 0
Link Bw CT5    : 0
Link Bw CT6    : 0
Link Bw CT7    : 0

Bandwidth Constraints for Class Types (Kbps)
BC0            : 0
BC1            : 0
BC2            : 0
BC3            : 0
BC4            : 0
BC5            : 0
BC6            : 0
BC7            : 0

Bandwidth for TE Class Types (Kbps)
TE0-> Resv. Bw : 0
Unresv. Bw     : 0
```

```

TE1-> Resv. Bw : 0          Unresv. Bw : 0
TE2-> Resv. Bw : 0          Unresv. Bw : 0
TE3-> Resv. Bw : 0          Unresv. Bw : 0
TE4-> Resv. Bw : 0          Unresv. Bw : 0
TE5-> Resv. Bw : 0          Unresv. Bw : 0
TE6-> Resv. Bw : 0          Unresv. Bw : 0
TE7-> Resv. Bw : 0          Unresv. Bw : 0
No Neighbors.
-----
Interface : ip-10.10.12.3
-----
Interface      : ip-10.10.12.3
Port ID       : 1/1/9
Admin State   : Up          Oper State      : Up
Active Sessions : 1         Active Resvs  : 0
Total Sessions : 1
Subscription  : 100 %       Port Speed    : 1000 Mbps
Total BW      : 1000 Mbps   Aggregate     : Dsabl
Hello Interval : 3000 ms    Hello Timeouts : 0
Authentication : Disabled
Auth Rx Seq Num : n/a       Auth Key Id   : n/a
Auth Tx Seq Num : n/a       Auth Win Size : n/a
Refresh Reduc.  : Disabled   Reliable Deli. : Disabled
Bfd Enabled     : No         Graceful Shut. : Disabled

Percent Link Bandwidth for Class Types
Link Bw CT0    : 100        Link Bw CT4    : 0
Link Bw CT1    : 0          Link Bw CT5    : 0
Link Bw CT2    : 0          Link Bw CT6    : 0
Link Bw CT3    : 0          Link Bw CT7    : 0

Bandwidth Constraints for Class Types (Kbps)
BC0            : 1000000     BC4            : 0
BC1            : 0          BC5            : 0
BC2            : 0          BC6            : 0
BC3            : 0          BC7            : 0

Bandwidth for TE Class Types (Kbps)
TE0-> Resv. Bw : 0          Unresv. Bw    : 1000000
TE1-> Resv. Bw : 0          Unresv. Bw    : 1000000
TE2-> Resv. Bw : 0          Unresv. Bw    : 1000000
TE3-> Resv. Bw : 0          Unresv. Bw    : 1000000
TE4-> Resv. Bw : 0          Unresv. Bw    : 1000000
TE5-> Resv. Bw : 0          Unresv. Bw    : 1000000
TE6-> Resv. Bw : 0          Unresv. Bw    : 1000000
TE7-> Resv. Bw : 0          Unresv. Bw    : 1000000
Neighbors      : 10.10.12.2
-----
Interface : ip-10.10.4.3
-----
Interface      : ip-10.10.4.3
Port ID       : 1/1/8
Admin State   : Up          Oper State      : Up
Active Sessions : 1         Active Resvs  : 0
Total Sessions : 1
Subscription  : 100 %       Port Speed    : 1000 Mbps
Total BW      : 1000 Mbps   Aggregate     : Dsabl
Hello Interval : 3000 ms    Hello Timeouts : 0
Authentication : Disabled
Auth Rx Seq Num : n/a       Auth Key Id   : n/a
Auth Tx Seq Num : n/a       Auth Win Size : n/a
Refresh Reduc.  : Disabled   Reliable Deli. : Disabled
Bfd Enabled     : No         Graceful Shut. : Disabled

```

```
Percent Link Bandwidth for Class Types
Link Bw CT0      : 100      Link Bw CT4      : 0
Link Bw CT1      : 0        Link Bw CT5      : 0
Link Bw CT2      : 0        Link Bw CT6      : 0
Link Bw CT3      : 0        Link Bw CT7      : 0

Bandwidth Constraints for Class Types (Kbps)
BC0              : 1000000   BC4              : 0
BC1              : 0        BC5              : 0
BC2              : 0        BC6              : 0
BC3              : 0        BC7              : 0

Bandwidth for TE Class Types (Kbps)
TE0-> Resv. Bw   : 0        Unresv. Bw      : 1000000
TE1-> Resv. Bw   : 0        Unresv. Bw      : 1000000
TE2-> Resv. Bw   : 0        Unresv. Bw      : 1000000
TE3-> Resv. Bw   : 0        Unresv. Bw      : 1000000
TE4-> Resv. Bw   : 0        Unresv. Bw      : 1000000
TE5-> Resv. Bw   : 0        Unresv. Bw      : 1000000
TE6-> Resv. Bw   : 0        Unresv. Bw      : 1000000
TE7-> Resv. Bw   : 0        Unresv. Bw      : 1000000
Neighbors       : 10.10.4.2
```

Interface : ip-10.10.2.3

```
Interface      : ip-10.10.2.3
Port ID       : 1/1/4
Admin State   : Up          Oper State      : Down
Active Sessions : 0        Active Resvs   : 0
Total Sessions : 0
Subscription  : 100 %      Port Speed     : 0 Mbps
Total BW      : 0 Mbps     Aggregate      : Dsabl
Hello Interval : 3000 ms   Hello Timeouts : 0
Authentication : Disabled
Auth Rx Seq Num : n/a      Auth Key Id    : n/a
Auth Tx Seq Num : n/a      Auth Win Size  : n/a
Refresh Reduc. : Disabled  Reliable Deli. : Disabled
Bfd Enabled    : No        Graceful Shut. : Disabled
```

```
Percent Link Bandwidth for Class Types
Link Bw CT0      : 100      Link Bw CT4      : 0
Link Bw CT1      : 0        Link Bw CT5      : 0
Link Bw CT2      : 0        Link Bw CT6      : 0
Link Bw CT3      : 0        Link Bw CT7      : 0
```

```
Bandwidth Constraints for Class Types (Kbps)
BC0              : 0        BC4              : 0
BC1              : 0        BC5              : 0
BC2              : 0        BC6              : 0
BC3              : 0        BC7              : 0
```

```
Bandwidth for TE Class Types (Kbps)
TE0-> Resv. Bw   : 0        Unresv. Bw      : 0
TE1-> Resv. Bw   : 0        Unresv. Bw      : 0
TE2-> Resv. Bw   : 0        Unresv. Bw      : 0
TE3-> Resv. Bw   : 0        Unresv. Bw      : 0
TE4-> Resv. Bw   : 0        Unresv. Bw      : 0
TE5-> Resv. Bw   : 0        Unresv. Bw      : 0
TE6-> Resv. Bw   : 0        Unresv. Bw      : 0
TE7-> Resv. Bw   : 0        Unresv. Bw      : 0
No Neighbors.
```

Table 20: Output fields: RSVP interface

Label	Description
Interface	Displays the name of the IP interface
Total Sessions	Displays the total number of RSVP sessions on this interface, including sessions that are active and sessions that have been signaled but a response has not yet been received
Active Sessions	Displays the total number of active RSVP sessions on this interface
Total BW (Mbps)	Displays the amount of bandwidth in megabits per second (Mbps) available to be reserved for the RSVP protocol on the interface
Resv BW (Mbps)	Displays the amount of bandwidth in mega-bits per seconds (Mbps) reserved on this interface. A value of zero (0) indicates that no bandwidth is reserved.
Adm	Down — the RSVP interface is administratively disabled Up — the RSVP interface is administratively enabled
Opr	Down — the RSVP interface is operationally down Up — the RSVP interface is operationally up
Port ID	Displays the physical port bound to the interface
Active Resvs	Displays the total number of active RSVP sessions that have reserved bandwidth
Subscription	Displays the percentage of the link bandwidth that RSVP can use for reservation. When the value is zero (0), no new sessions are permitted on this interface.
Port Speed	Displays the speed for the interface
Unreserved BW	Displays the amount of unreserved bandwidth
Reserved BW	Displays the amount of bandwidth, in megabits per second (Mbps), reserved by the RSVP session on this interface. A value of zero (0) indicates that no bandwidth is reserved.
Total BW	Displays the amount of bandwidth, in megabits per second (Mbps), available to be reserved for the RSVP protocol on this interface
Hello Interval	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. When the value is zero (0), the sending of hello messages is disabled.

Label	Description
Refresh Time	Displays the interval between the successive Path and Resv refresh messages. RSVP declares the session down after it misses $((\text{keep-multiplier} + 0.5) * 1.5 * \text{refresh-time})$ consecutive refresh messages.
Hello Timeouts	Displays the total number of hello messages that timed out on this RSVP interface
Neighbors	Displays the IP address of the RSVP neighbor
Sent	Displays the total number of error free RSVP packets that have been transmitted on the RSVP interface
Recd	Displays the total number of error free RSVP packets received on the RSVP interface
Total Packets	Displays the total number of RSVP packets, including errors, received on the RSVP interface
Bad Packets	Displays the total number of RSVP packets with errors transmitted on the RSVP interface
Paths	Displays the total number of RSVP PATH messages received on the RSVP interface
Path Errors	Displays the total number of RSVP PATH ERROR messages transmitted on the RSVP interface
Path Tears	Displays the total number of RSVP PATH TEAR messages received on the RSVP interface
Resvs	Displays the total number of RSVP RESV messages received on the RSVP interface
Resv Confirms	Displays the total number of RSVP RESV CONFIRM messages received on the RSVP interface
Resv Errors	Displays the total number of RSVP RESV ERROR messages received on the RSVP interface
Resv Tears	Displays the total number of RSVP RESV TEAR messages received on the RSVP interface
Refresh Summaries	Displays the total number of RSVP RESV summary refresh messages received on the interface
Refresh Acks	Displays the total number of RSVP RESV acknowledgment messages received when refresh reduction is enabled on the RSVP interface
Hellos	Displays the total number of RSVP RESV HELLO REQ messages received on the interface

Label	Description
Bfd Enabled	Yes — BFD is enabled on the RSVP interface No — BFD is disabled on the RSVP interface

neighbor

Syntax

neighbor [*ip-address*] [*detail*]

Context

show>router>rsvp

Platforms

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

Description

This command displays neighbor information.

Parameters

ip-address

Displays RSVP information about the specified IP address.

detail

Keyword to display detailed information.

Output

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

Sample output

```
*A:Dut>config>router>mpls# show router rsvp neighbor
=====
RSVP Neighbors
=====
Legend :
  LR - Local Refresh Reduction      RR - Remote Refresh Reduction
  LD - Local Reliable Delivery      RM - Remote Node supports
Message ID
  LG - Local Graceful Restart      RG - Remote Graceful Restart
=====
Neighbor      Interface      Hello  Last Oper
Flags                                     Change
=====
10.20.1.2     ip-10.10.1.1      N/A   0d 00:00:44
10.20.1.3     ip-10.10.2.1      N/A   0d 00:00:44
-----
Neighbors : 2
```

*A:Dut>config>router>mpls#

Table 21: Output fields: RSVP neighbor

Label	Description
Neighbor	Displays the IP address of the RSVP neighbor
Interface	Displays the interface ID of the RSVP neighbor
Hello	Displays the status of the Hello message
Last Oper Change	Displays the time of the last operational change to the connection
Flags	Displays the flags that are associated with the connection to the neighbor

session

Syntax

session *session-type* [**from** *ip-address* | **to** *ip-address*] [**lsp-name** *name*] [**status** {**up** | **down**}] [**detail**]

Context

show>router>rsvp

Platforms

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

Description

This command shows RSVP session information.

Parameters

session *session-type*

Specifies the session type.

Values originate, transit, terminate, detour, detour-transit, detour-terminate, bypass-tunnel, manual-bypass, s2l-originate, s2l-transit, s2l-terminate, s2l-bypass-tunnel

from *ip-address*

Specifies the IP address of the originating router.

to *ip-address*

Specifies the IP address of the egress router.

lsp-name name

Specifies the name of the LSP used in the path.

up

Keyword to display a session that is operationally up.

down

Keyword to display a session that is operationally down.

detail

Keyword to display detailed information.

Output

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

Sample output

```
*A:SRU4>show>router>rsvp# session
=====
RSVP Sessions
=====
From          To          Tunnel LSP   Name                               State
ID            ID
-----
10.20.1.5     10.20.1.4   18      27648 b4-1::b4-1                        Up
10.20.1.5     10.20.1.4   1       37902 gsr::gsr                          Up
10.20.1.5     10.20.1.22  11      53760 to_10_20_1_22_cspf::to_10_2* Up
10.20.1.4     10.20.1.20  146     17920 to_10_20_1_20_cspf_3::to_10* Up
10.20.1.4     10.20.1.20  145     34816 to_10_20_1_20_cspf_2::to_10* Up
10.20.1.4     10.20.1.20  147     45056 to_10_20_1_20_cspf_4::to_10* Up
10.20.1.4     10.20.1.20  148     6656  to_10_20_1_20_cspf_5::to_10* Up
10.20.1.4     10.20.1.20  149     58880 to_10_20_1_20_cspf_6::to_10* Up
10.20.1.4     10.20.1.20  150     13312 to_10_20_1_20_cspf_7::to_10* Up
10.20.1.4     10.20.1.20  152     40448 to_10_20_1_20_cspf_9::to_10* Up
10.20.1.4     10.20.1.20  154     27648 to_10_20_1_20_cspf_11::to_1* Up
10.20.1.4     10.20.1.20  155     12288 to_10_20_1_20_cspf_12::to_1* Up
10.20.1.4     10.20.1.20  151     46080 to_10_20_1_20_cspf_8::to_10* Up
10.20.1.4     10.20.1.20  153     512   to_10_20_1_20_cspf_10::to_1* Up
10.20.1.4     10.20.1.22  164     62464 to_10_20_1_22_cspf_2::to_10* Up
10.20.1.4     10.20.1.20  156     37888 to_10_20_1_20_cspf_13::to_1* Up
10.20.1.4     10.20.1.20  157     24064 to_10_20_1_20_cspf_14::to_1* Up
10.20.1.4     10.20.1.20  158     19968 to_10_20_1_20_cspf_15::to_1* Up
10.20.1.4     10.20.1.20  161     59904 to_10_20_1_20_cspf_18::to_1* Up
...
10.20.1.3     10.20.1.4   54      23088 to_110_20_1_4_cspf_4::to_11* Up
-----
Sessions : 1976
=====
* indicates that the corresponding row element may have been truncated.
*A:SRU4>show>router>rsvp#

A:ALA-12# show router rsvp session lsp-name A_C_2::A_C_2 status up
=====
RSVP Sessions
=====
From          To          Tunnel LSP   Name                               State
ID            ID
-----
10.20.1.1     10.20.1.3   2       40    A_C_2::A_C_2                      Up
```

```
-----  
Sessions : 1  
=====
```

A:ALA-12#

Table 22: Output fields: RSVP session

Label	Description
From	Displays the IP address of the originating router
To	Displays the IP address of the egress router
Tunnel ID	Displays the IP address of the tunnel ingress node supporting this RSVP session
LSP ID	Displays the ID assigned by the agent to this RSVP session
Name	Displays the administrative name assigned to the RSVP session by the agent
State	Down — the operational state of this RSVP session is down
	Up — the operational state of this RSVP session is up

statistics

Syntax
statistics

Context
show>router>rsvp

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

Description
This command displays global statistics in the RSVP instance.

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

Sample output

```
*A:SRU4>show>router>rsvp# statistics  
=====
```

RSVP Global Statistics

```
=====
```

```
PATH Timeouts      : 1026          RESV Timeouts      : 182
=====
*A:SRU4>show>router>rsvp#
```

Table 23: Output fields: RSVP statistics

Label	Description
PATH Timeouts	Displays the total number of path timeouts
RESV Timeouts	Displays the total number of RESV timeouts

status

Syntax

rsvp status

Context

show>router>rsvp

Platforms

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

Description

This command displays RSVP status.

Output

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

Sample output

```
*A:SRU4>show>router>rsvp# status
=====
RSVP Status
=====
Admin Status      : Up          Oper Status      : Up
Keep Multiplier   : 3           Refresh Time     : 30 sec
Message Pacing    : Disabled    Pacing Period    : 100 msec
Max Packet Burst  : 650 msg     Refresh Bypass   : Disabled
=====
*A:SRU4>show>router>rsvp#
```

Table 24: Output fields: RSVP status

Label	Description
Admin Status	Down — RSVP is administratively disabled

Label	Description
	Up — RSVP is administratively enabled
Oper Status	Down — RSVP is operationally down Up — RSVP is operationally up
Keep Multiplier	Displays the keep-multiplier <i>number</i> used by RSVP to declare that a reservation is down or the neighbor is down.
Refresh Time	Displays the refresh-time interval, in seconds, between the successive Path and Resv refresh messages
Message Pacing	Enabled — RSVP messages, specified in the max-burst command, are sent in a configured interval, specified in the period command Disabled — Message pacing is disabled; RSVP message transmission is not regulated
Pacing Period	Displays the time interval, in milliseconds, when the router can send the specified number of RSVP messages specified in the rsvp max-burst command
Max Packet Burst	Displays the maximum number of RSVP messages that are sent in the specified period under normal operating conditions

2.17.2.4 Tools commands

cspf

Syntax

cspf *to ip-addr* [**from** *ip-addr*] [**bandwidth** *bandwidth*] [**include-bitmap** *bitmap*] [**exclude-bitmap** *bitmap*] [**hop-limit** *limit*] [**exclude-address** *excl-addr* [*excl-addr...*(up to 8 max)]] [**use-te-metric**] [**strict-srlg**] [**srlggroup** *grp-id...*(up to 8 max)] [**skip-interface** *interface-name*]

Context

tools>perform>router>mpls

Platforms

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

Description

This command computes a CSPF path with specified user constraints.

Parameters

to *ip-addr*

Specifies the destination IP address.

from *ip-addr*

Specifies the originating IP address.

bandwidth *bandwidth*

Specifies the amount of bandwidth in mega-bits per second (Mbps) to be reserved.

include-bitmap *bitmap*

Specifies to include a bit-map that specifies a list of admin groups that should be included during setup.

exclude-bitmap *bitmap*

Specifies to exclude a bit-map that specifies a list of admin groups that should be included during setup.

hop-limit *limit*

Specifies the total number of hops a detour LSP can take before merging back onto the main LSP path.

exclude-address *ip-addr*

Specifies IP addresses, up to 8, that should be included during setup.

use-te-metric

Keyword to specify the use of the traffic engineering metric used on the interface.

strict-srlg

Keyword to specify whether to associate the LSP with a bypass or signal a detour if a bypass or detour satisfies all other constraints except the SRLG constraints.

srlg-group *grp-id*

Specifies up to 8 Shared Risk Link Groups (SRLGs). An SRLG group represents a set of interfaces which could be subject to the same failures or defects and therefore, share the same risk of failing.

Values 0 to 4294967295

skip-interface *interface-name*

Specifies an interface name that should be skipped during setup.

Output

Sample output

```
*A:Dut-C# tools perform router mpls cspf to 10.20.1.6

Req CSPF for all ECMP paths
  from: this node to: 10.20.1.6 w/
(no Diffserv) class: 0 , setup Priority 7, Hold Priority 0 TE Class: 7

CSPF Path
To       : 10.20.1.6
Path 1   : (cost 2000)
Src:     10.20.1.3 (= Rtr)
```



```
Egr:    unnumbered lnkId 4      -
> Ingr:  unnumbered lnkId 2      Rtr:   10.20.1.5      (met 1000)
Egr:    unnumbered lnkId 3      -
> Ingr:  unnumbered lnkId 3      Rtr:   10.20.1.6      (met 1000)
  Dst:   10.20.1.6    (= Rtr)

Path 2   : (cost 2000)
  Src:   10.20.1.3    (= Rtr)
Egr:    unnumbered lnkId 5      -
> Ingr:  unnumbered lnkId 5      Rtr:   10.20.1.4      (met 1000)
Egr:    unnumbered lnkId 3      -
> Ingr:  unnumbered lnkId 2      Rtr:   10.20.1.6      (met 1000)
  Dst:   10.20.1.6    (= Rtr)

*A:Dut-C#
```

resignal

Syntax
resignal {*lsp lsp-name path path-name* | **delay** *minutes*}

Context
tools>perform>router>mpls

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

Description
This command resignals specified LSP paths. The *minutes* parameter configures the global timer to resignal all LSPs. If only *lsp-name* and *path-name* are provided, the specified LSP is resigned immediately.

- Parameters**
- lsp-name***
Specifies a unique LSP name, up to 32 characters.
 - path-name***
Specifies the name for the LSP path, up to 32 characters.
 - minutes***
Specifies the delay interval, in minutes, before all LSPs are resigned. If the value 0 is entered, all LSPs are resigned immediately.
- Values** 0 to 30

resignal-bypass

Syntax

resignal-bypass {*lsp bypass-lsp-name* [**force**] | **delay** *minutes*}

Context

tools>perform>router>mpls

Platforms

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

Description

This command performs a manual re-optimization of a specific dynamic or manual bypass LSP, or of all dynamic bypass LSPs.

The user must configure the manual bypass LSP name. The dynamic bypass LSP name is shown in the output of **show>router>mpls>bypass-tunnel dynamic detail**.

The **delay** option triggers the global reoptimization of all dynamic bypass LSPs at the expiry of the specified delay. This option forces the global bypass resignal timer to expire after an amount of time equal to the value of the **delay** parameter. This option has no effect on a manual bypass LSP.

However, when *bypass-lsp-name* is specified, the named dynamic or manual bypass LSP is signaled, and the associations are moved only if the new bypass LSP path has a lower cost than the current path. This behavior is different from that of the **tools>perform>router>mpls>resignal** command for the primary or secondary active path of an LSP, which signals and switches to the new path, regardless of the cost comparison. This handling is required because a bypass LSP may have a large number of PSB associations and the processing churn is much higher.

In the specific case where the name corresponds to a manual bypass LSP, the LSP is torn down and resignaled using the new path provided by CSPF if no PSB associations exist. If one or more PSB associations exist but no PLR is active, the command fails and the user is required to explicitly enter the **force** option. In this case, the manual bypass LSP is torn down and resignaled, temporarily leaving the associated LSP primary paths unprotected. If one or more PLRs associated with the manual bypass LSP is active, the command fails.

Finally, and as with the timer-based resignal, the PSB associations are checked for the SRLG and admin-group constraints using the updated information provided by CSPF for the current path and new path of the bypass LSP.

Parameters

***lsp bypass-lsp-name* [**force**]**

Specifies the name, up to 160 characters, of the dynamic or manual bypass LSP. The **force** option is required when the name corresponds to a manual bypass LSP and the LSP has PSB associations.

delay minutes

Specifies the time, in minutes, that MPLS waits before attempting to resignal dynamic bypass LSP paths originated on the system.

Values 0 to 30

tp-tunnel

Syntax

tp-tunnel

Context

tools>perform>router>mpls

Platforms

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

Description

Commands in this context perform linear protection operations on an MPLS-TP LSP.

clear

Syntax

clear id *tunnel-id* *lsp-name*

Context

tools>perform>router>mpls>tp-tunnel

Platforms

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

Description

This command clears all MPLS-TP linear protection operational commands for the specified LSPs that are currently active.

Parameters

tunnel-id

Specifies the tunnel number of the MPLS-TP LSP, up to 32 characters.

lsp-name

Specifies the name of the MPLS-TP LSP.

Values 1 to 61440

force

Syntax

force id *tunnel-id lsp-name*

Context

tools>perform>router>mpls>tp-tunnel

Platforms

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

Description

This command performs a force switchover of the MPLS-TP LSP from the active path to the protect path.

Parameters

tunnel-id

Specifies the tunnel number of the MPLS-TP LSP, up to 32 characters.

lsp-name

Specifies the name of the MPLS-TP LSP.

Values 1 to 61440

lockout

Syntax

lockout *tunnel-id lsp-name*

Context

tools>perform>router>mpls>tp-tunnel

Platforms

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

Description

This command performs a lockout of protection for an MPLS-TP LSP. This prevents a switchover to the protect path.

Parameters

tunnel-id

Specifies the tunnel number of the MPLS-TP LSP, up to 32 characters.

lsp-name

Specifies the name of the MPLS-TP LSP.

Values 1 to 61440

manual

Syntax

manual *tunnel-id lsp-name*

Context

tools>perform>router>mpls>tp-tunnel

Platforms

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

Description

This command performs a manual switchover of the MPLS-TP LSP from the active path to the protect path.

Parameters

tunnel-id

Specifies the tunnel number of the MPLS-TP LSP, up to 32 characters.

lsp-name

Specifies the name of the MPLS-TP LSP.

Values 1 to 61440

trap-suppress

Syntax

trap-suppress *number-of-traps time-interval*

Context

tools>perform>router>mpls

Platforms

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

Description

This command modifies thresholds for trap suppression. The command is used to suppress traps after a specified number of traps has been raised within the specified period of time.

Parameters

number-of-traps

Specifies the number of traps, in multiples of 100.

Values 100 to 1000

time-interval

Specifies the time interval, in seconds.

Values 1 to 300

2.17.2.5 Clear commands

fec-egress-statistics

Syntax

fec-egress-statistics [*ip-prefix/mask*]

Context

clear>router>ldp

Platforms

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

Description

This command resets or clears LDP FEC egress statistics.

Parameters

ip-prefix

Specifies information for the specified IP prefix and mask length. Host bits must be "0".

mask

Specifies the 32-bit address mask used to indicate the bits of an IP address that are being used for the subnet address.

Values 0 to 32

interface

Syntax

interface *ip-int-name*

Context

clear>router>mpls

Platforms

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

Description

This command resets or clears statistics for MPLS interfaces.

Parameters

ip-int-name

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

lsp

Syntax

lsp *lsp-name*

Context

clear>router>mpls

Platforms

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

Description

This command resets and restarts an LSP.

Parameters

lsp-name

Specifies the name of the LSP to clear, up to 64 characters.

interface

Syntax

interface *ip-int-name* **statistics**

Context

clear>router>rsvp

Platforms

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

Description

This command resets or clears statistics for an RSVP interface.

Parameters

ip-int-name

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

statistics

Keyword to clear only statistics.

statistics

Syntax

statistics

Context

clear>router>rsvp

Platforms

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

Description

This command clears global statistics for the RSVP instance, for example, clears **path** and **resv timeout** counters.

2.17.2.6 Debug commands

mpls

Syntax

mpls [**lsp** *lsp-name*] [**sender** *source-address*] [**endpoint** *endpoint-address*] [**tunnel-id** *tunnel-id*] [**lsp-id** *lsp-id*]

no mpls

Context

debug>router

Platforms

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

Description

This command enables and configures debugging for MPLS.

Parameters

lsp-name

Specifies the name that identifies the LSP, up to 32 characters.

source-address

Specifies the system IP address of the sender.

endpoint-address

Specifies the far-end system IP address.

tunnel-id

Specifies the MPLS SDP ID.

Values 0 to 4294967295

lsp-id

Specifies the LSP ID.

Values 1 to 65535

event

Syntax

[no] event

Context

debug>router>mpls

debug>router>rsvp

Platforms

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

Description

This command enables debugging for specific events.

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

all

Syntax

all [detail]

no all

Context

debug>router>mpls>event

debug>router>rsvp>event

Platforms

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

Description

This command debugs all events.

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

Parameters

detail

Keyword to display detailed information about all events.

auth

Syntax

auth

no auth

Context

debug>router>rsvp>event

Platforms

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

Description

This command debugs authentication events.

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

Parameters

detail

Keyword to display detailed information about authentication events.

frr

Syntax

frr [detail]

no frr

Context

debug>router>mpls>event

Platforms

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

Description

This command debugs fast re-route events.

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

Parameters

detail

Keyword to display detailed information about re-route events.

iom

Syntax

iom [detail]

no iom

Context

debug>router>mpls>event

Platforms

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

Description

This command debugs MPLS IOM events.

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

Parameters

detail

Keyword to display detailed information about MPLS IOM events.

lsp-setup

Syntax

lsp-setup [detail]

no lsp-setup

Context

debug>router>mpls>event

Platforms

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

Description

This command debugs LSP setup events.

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

Parameters

detail

Keyword to display detailed information about LSP setup events.

mbb

Syntax

mbb [detail]

no mbb

Context

debug>router>mpls>event

Platforms

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

Description

This command debugs the state of the most recent invocation of the make-before-break (MBB) functionality.

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

Parameters

detail

Keyword to display detailed information about MBB events.

misc

Syntax

misc [detail]

no misc

Context

debug>router>mpls>event

debug>router>rsvp>event

Platforms

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

Description

This command debugs miscellaneous events.

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

Parameters

detail

Keyword to display detailed information about miscellaneous events.

XC

Syntax

xc [**detail**]

no xc

Context

debug>router>mpls>event

Platforms

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

Description

This command debugs cross connect events.

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

Parameters

detail

Keyword to display detailed information about cross connect events.

rsvp

Syntax

[**lsp** *lsp-name*] [**sender** *source-address*] [**endpoint** *endpoint-address*] [**tunnel-id** *tunnel-id*] [**lsp-id** *lsp-id*]
[**interface** *ip-int-name*]

no rsvp

Context

debug>router

Platforms

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

Description

This command enables and configures debugging for RSVP.

Parameters

lsp-name

Specifies the name that identifies the LSP, up to 32 characters.

sender source-address

Displays the system IP address of the sender.

endpoint-address

Displays the far-end system IP address.

tunnel-id

Displays the RSVP tunnel ID.

Values 0 to 4294967295

lsp-id

Displays the LSP ID.

Values 1 to 65535

ip-int-name

Displays the interface name. The interface name can be up to 32 characters long and must be unique. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

nbr

Syntax

nbr [detail]

no nbr

Context

debug>router>rsvp>event

Platforms

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

Description

This command debugs neighbor events.

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

Parameters

detail

Keyword to display detailed information about neighbor events.

path

Syntax

path [detail]

no path

Context

debug>router>rsvp>event

Platforms

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

Description

This command debugs path-related events.

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

Parameters

detail

Keyword to display detailed information about path-related events.

resv

Syntax

resv [detail]

no resv

Context

debug>router>rsvp>event

Platforms

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

Description

This command debugs RSVP reservation events.

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

Parameters

detail

Keyword to display detailed information about RSVP reservation events.

```
rr
```

Syntax

```
rr
```

```
no rr
```

Context

```
debug>router>rsvp>event
```

Platforms

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

Description

This command debugs refresh reduction events.

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

Parameters

detail

Keyword to display detailed information about refresh reduction events.

```
packet
```

Syntax

```
[no] packet
```

Context

```
debug>router>rsvp>
```

Platforms

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

Description

This command enable the context to debug packets.

ack

Syntax

ack [detail]

no ack

Context

debug>router>rsvp>packet

Platforms

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

Description

This command debugs ACK packets.

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

Parameters

detail

Keyword to display detailed information about RSVP-TE ACK packets.

bundle

Syntax

bundle [detail]

no bundle

Context

debug>router>rsvp>packet

Platforms

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

Description

This command debugs bundle packets.

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

Parameters

detail

Keyword to display detailed information about RSVP-TE bundle packets.

all

Syntax

all [detail]

no all

Context

debug>router>rsvp>packet

Platforms

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

Description

This command debugs all packets.

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

Parameters

detail

Keyword to display detailed information about all RSVP packets.

hello

Syntax

hello [detail]

no hello

Context

debug>router>rsvp>packet

Platforms

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

Description

This command debugs hello packets.

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

Parameters

detail

Keyword to display detailed information about hello packets.

path

Syntax

path [**detail**]

no path

Context

debug>router>rsvp>packet

Platforms

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

Description

This command enables debugging for RSVP path packets.

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

Parameters

detail

Keyword to display detailed information about path-related events.

patherr

Syntax

patherr [**detail**]

no patherr

Context

debug>router>rsvp>packet

Platforms

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

Description

This command debugs path error packets.

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

Parameters

detail

Keyword to display detailed information about path error packets.

pathtear

Syntax

pathtear [detail]

no pathtear

Context

debug>router>rsvp>packet

Platforms

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

Description

This command debugs path tear packets.

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

Parameters

detail

Keyword to display detailed information about path tear packets.

resv

Syntax

resv [detail]

no resv

Context

debug>router>rsvp>packet

Platforms

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

Description

This command enables debugging for RSVP resv packets.

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

Parameters

detail

Keyword to display detailed information about RSVP Resv events.

resverr

Syntax

resverr [detail]

no resverr

Context

debug>router>rsvp>packet

Platforms

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

Description

This command debugs ResvErr packets.

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

Parameters

detail

Keyword to display detailed information about ResvErr packets.

resvtear

Syntax

resvtear [detail]

no resvtear

Context

debug>router>rsvp>packet

Platforms

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

Description

This command debugs ResvTear packets.

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

Parameters

detail

Keyword to display detailed information about ResvTear packets.

srefresh

Syntax

srefresh [detail]

no srefresh

Context

debug>router>rsvp>packet

Platforms

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

Description

This command debugs srefresh packets.

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

Parameters

detail

Keyword to display detailed information about RSVP-TE srefresh packets.

3 Label Distribution Protocol

This chapter provides information to enable Label Distribution Protocol (LDP).

3.1 Label Distribution Protocol

Label Distribution Protocol (LDP) is a protocol used to distribute labels in non-traffic-engineered applications. LDP allows routers to establish label switched paths (LSPs) through a network by mapping network-layer routing information directly to data link layer-switched paths.

An LSP is defined by the set of labels from the ingress Label Switching Router (LSR) to the egress LSR. LDP associates a Forwarding Equivalence Class (FEC) with each LSP it creates. A FEC is a collection of common actions associated with a class of packets. When an LSR assigns a label to a FEC, it must allow other LSRs in the path know about the label. LDP helps to establish the LSP by providing a set of procedures that LSRs can use to distribute labels.

The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each LSR splices incoming labels for a FEC to the outgoing label assigned to the next hop for the specified FEC.

LDP allows an LSR to request a label from a downstream LSR so it can bind the label to a specific FEC. The downstream LSR responds to the request from the upstream LSR by sending the requested label.

LSRs can distribute a FEC label binding in response to an explicit request from another LSR. This is known as Downstream On Demand (DOD) label distribution. LSRs can also distribute label bindings to LSRs that have not explicitly requested them. This is called Downstream Unsolicited (DUS).

3.1.1 LDP and MPLS

LDP performs the label distribution only in MPLS environments. The LDP operation begins with a hello discovery process to find LDP peers in the network. LDP peers are two LSRs that use LDP to exchange label/FEC mapping information. An LDP session is created between LDP peers. A single LDP session allows each peer to learn the other's label mappings (LDP is bidirectional) and to exchange label binding information.

LDP signaling works with the MPLS label manager to manage the relationships between labels and the corresponding FEC. For service-based FECs, LDP works in tandem with the Service Manager to identify the virtual leased lines (VLLs) and Virtual Private LAN Services (VPLSs) to signal.

An MPLS label identifies a set of actions that the forwarding plane performs on an incoming packet before discarding it. The FEC is identified through the signaling protocol (in this case, LDP) and allocated a label. The mapping between the label and the FEC is communicated to the forwarding plane. In order for this processing on the packet to occur at high speeds, optimized tables are maintained in the forwarding plane that enable fast access and packet identification.

When an unlabeled packet ingresses the IP/MPLS router, classification policies associate it with a FEC. The appropriate label is imposed on the packet, and the packet is forwarded. Other actions that can take place before a packet is forwarded are imposing additional labels, other encapsulations, learning actions, and so on. When all actions associated with the packet are completed, the packet is forwarded.

When a labeled packet ingresses the router, the label or stack of labels indicates the set of actions associated with the FEC for that label or label stack. The actions are preformed on the packet and then the packet is forwarded.

The LDP implementation provides DOD, DUS, ordered control, liberal label retention mode support.

3.1.2 LDP architecture

LDP comprises a few processes that handle the protocol PDU transmission, timer-related issues, and protocol state machine. The number of processes is kept to a minimum to simplify the architecture and to allow for scalability. Scheduling within each process prevents starvation of any particular LDP session, while buffering alleviates TCP-related congestion issues.

The LDP subsystems and their relationships to other subsystems are illustrated in [Figure 29: Subsystem interrelationships](#). This illustration shows the interaction of the LDP subsystem with other subsystems, including memory management, label management, service management, SNMP, interface management, and RTM. In addition, debugging capabilities are provided through the logger.

Communication within LDP tasks is typically done by inter-process communication through the event queue, as well as through updates to the various data structures. The primary data structures that LDP maintains are:

- **FEC/label database**

This database contains all the FEC to label mappings that include, both sent and received. It also contains both address FECs (prefixes and host addresses) as well as service FECs (L2 VLLs and VPLS).

- **timer database**

This database contains all the timers for maintaining sessions and adjacencies.


- **session database**

This database contains all the session and adjacency records, and serves as a repository for the LDP MIB objects.

3.1.3 Subsystem interrelationships

The following figure shows how LDP and the other subsystems work to provide services.

Figure 29: Subsystem interrelationships



3.1.3.1 Memory manager and LDP

LDP does not use any memory until it is instantiated. It preallocates some amount of fixed memory so that initial startup actions can be performed. Memory allocation for LDP comes out of a pool reserved for LDP that can grow dynamically as needed. Fragmentation is minimized by allocating memory in larger chunks and managing the memory internally to LDP. When LDP is shut down, it releases all memory allocated to it.

3.1.3.2 Label manager

LDP assumes that the label manager is up and running. LDP will abort initialization if the label manager is not running. The label manager is initialized at system boot-up; therefore, anything that causes it to fail will likely imply that the system is not functional. The 7210 devices use a label range from 28672 (28K) to 131071 (128K-1) to allocate all dynamic labels, including RSVP allocated labels and VC labels.

3.1.3.3 LDP configuration

The 7210 SAS devices use a single consistent interface to configure all protocols and services. CLI commands are translated to SNMP requests and are handled through an agent-LDP interface. LDP can be instantiated or deleted through SNMP. Also, LDP targeted sessions can be set up to specific endpoints. Targeted-session parameters are configurable.

3.1.3.4 Logger

LDP uses the logger interface to generate debug information relating to session setup and teardown, LDP events, label exchanges, and packet dumps. Per-session tracing can be performed.

3.1.3.5 Service manager

All interaction occurs between LDP and the service manager, because LDP is used primarily to exchange labels for Layer 2 services. In this context, the service manager informs LDP when an LDP session is to be set up or torn down, and when labels are to be exchanged or withdrawn. In turn, LDP informs service manager of relevant LDP events, such as connection setups and failures, timeouts, labels signaled/withdrawn.

3.1.4 Execution flow

LDP activity is limited to service-related signaling. Therefore, the configurable parameters are restricted to system-wide parameters, such as hello and keepalive timeouts.

3.1.4.1 Initialization

MPLS must be enabled when LDP is initialized. LDP makes sure that the various prerequisites, such as ensuring the system IP interface is operational, the label manager is operational, and there is memory available, are met. It then allocates itself a pool of memory and initializes its databases.

3.1.4.2 Session lifetime

In order for a targeted LDP (T-LDP) session to be established, an adjacency must be created. The LDP extended discovery mechanism requires hello messages to be exchanged between two peers for session establishment. After the adjacency establishment, session setup is attempted.

3.1.4.2.1 Session establishment

When the LDP adjacency is established, the session setup follows as per the LDP specification. Initialization and keepalive messages complete the session setup, followed by address messages to exchange all interface IP addresses. Periodic keepalives or other session messages maintain the session liveliness.

Because TCP is back-pressured by the receiver, it is necessary to be able to push that back-pressure all the way into the protocol. Packets that cannot be sent are buffered on the session object and re-attempted as the back-pressure eases.

3.1.5 Label exchange

Label exchange is initiated by the service manager. When an SDP is attached to a service (for example, the service gets a transport tunnel), a message is sent from the service manager to LDP. This causes a label mapping message to be sent. Additionally, when the SDP binding is removed from the service, the VC label is withdrawn. The peer must send a label release to confirm that the label is not in use.

3.1.5.1 Other reasons for label actions

Other reasons for label actions include:

- **MTU changes**

LDP withdraws the previously assigned label, and re-signals the FEC with the new MTU in the interface parameter.

- **Clear labels**

When a service manager command is issued to clear the labels, the labels are withdrawn, and new label mappings are issued.

- **SDP down**

When an SDP goes administratively down, the VC label associated with that SDP for each service is withdrawn.

- **Memory allocation failure**

If there is no memory to store a received label, it is released.

- **VC type unsupported**

When an unsupported VC type is received, the received label is released.

3.1.5.2 Cleanup

LDP closes all sockets, frees all memory, and shuts down all its tasks when it is deleted, so its memory usage is 0 when it is not running.

3.1.5.3 Configuring implicit null label

The implicit null label option allows a 7210 SAS egress LER to receive MPLS packets from the previous hop without the outer LSP label. The operation of the previous hop is referred to as penultimate hop popping (PHP). This option is signaled by the egress LER to the previous hop during the FEC signaling by the control protocol.

The user can configure to signal the implicit null option for all RSVP FECs for which this node is the egress LER using the following command:

```
config>router>rsvp>implicit-null-label
```

When the user changes the implicit null configuration option, RSVP withdraws all the FECs and re-advertises them using the new label value.

3.1.5.4 LDP filters

Both inbound and outbound LDP label binding filtering is supported.

Inbound filtering (import policy) allows configuration of a policy to control the label bindings an LSR accepts from its peers. Label bindings can be filtered based on:

- Neighbor: Match on bindings received from the specified peer
- Prefix-list: Match on bindings with the specified prefix/prefixes



Note:

The default import behavior is to accept all FECs received from peers. The LDP export policy can be used to explicitly add FECs (or non-LDP routes) for label propagation and does not filter out or stop propagation of any FEC received from neighbors.

Export policy enables configuration of a policy to advertise label bindings based on:

- Direct: All local subnets
- Prefix-list: Match on bindings with the specified prefix or prefixes



Note:

The LDP export policy will not filter out FECs. It is only used to explicitly add FECs (or non-LDP routes) for label propagation.

The default export behavior originates label bindings for system address and propagate all FECs received.

3.1.6 ECMP support for LDP



Note:

- LDP LER ECMP is not supported.
- LDP LSR ECMP is only supported on 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE.

This feature performs load balancing for LDP-based LSPs by having multiple outgoing next-hops for a specified IP prefix on ingress and transit LSRs.

An LSR that has multiple equal cost paths to a specified IP prefix can receive an LDP label mapping for this prefix from each downstream next-hop peer. As the LDP implementation uses the liberal label retention mode to retain all the labels for an IP prefix received from multiple next-hop peers.

Without ECMP support (only for LDP LSR LSPs on 7210 SAS), only one of these next-hop peers will be selected and installed in the forwarding plane. The next-hop peer selection algorithm looks up the route information obtained from the RTM for this prefix and finds the first valid LDP next-hop peer (for example, the first neighbor in the RTM entry from which a label mapping was received). If, for some reason, the outgoing label to the installed next-hop is no longer valid (for example, if the session to the peer is lost or the peer withdraws the label) a new valid LDP next-hop peer will be selected out of the existing next-hop peers and LDP will reprogram the forwarding plane to use the label sent by this peer.

With ECMP support, all the valid LDP next-hop peers, those that sent a label mapping for a specified IP prefix, will be installed in the forwarding plane. In transit LSR, an ingress label will be mapped to the next hops that are in the RTM and from which a valid mapping label has been received. The forwarding plane will then use an internal hashing algorithm to determine how the traffic will be distributed amongst these multiple next-hops, assigning each “flow” to a particular next-hop.

For more information about the hash algorithms at transit LSR, see “LAG and ECMP Hashing” in the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide*.

3.1.6.1 Label operations

If an LSR is the ingress for a specific IP prefix, LDP programs a push operation for the prefix in the forwarding engine. This creates an LSP ID to the Next Hop Label Forwarding Entry (NHLFE) (LTN) mapping and an LDP tunnel entry in the forwarding plane. LDP will also inform the Tunnel Table Manager (TTM) of this tunnel. Both the LTN entry and the tunnel entry will have a NHLFE for the label mapping that the LSR received from each of its next-hop peers.

If the LSR is to behave as a transit for a specified IP prefix, LDP will program a swap operation for the prefix in the forwarding engine. This results in the creation of an Incoming Label Map (ILM) entry in the forwarding plane. The ILM entry will have to map an incoming label to possibly multiple NHLFEs. If the LSR is an egress for a specific IP prefix, LDP programs a POP entry in the forwarding engine. Programming a POP entry results in an ILM entry in the forwarding plane, but with no NHLFEs.

When unlabeled packets arrive at the ingress LER, the forwarding plane will consult the LTN entry and will use a hashing algorithm to map the packet to one of the NHLFEs (push label) and forward the packet to the corresponding next-hop peer. For labeled packets arriving at a transit or egress LSR, the forwarding plane will consult the ILM entry and either use a hashing algorithm to map it to one of the NHLFEs if they exist (swap label) or simply route the packet if there are no NHLFEs (pop label).

Static FEC swap will not be activated unless there is a matching route in system route table that also matches the user configured static FEC next-hop.

3.1.6.2 LDP LSR ECMP hashing

The following table lists the cases in which LDP LSR ECMP hashing occurs when an MPLS encapsulated packet is received at LSR, and the cases where the MAC or IP packet address fields that are used in hashing vary.

Table 25: LSR hashing scenarios

Number and types of labels egressing iLER	Packet header address fields used in hashing ¹⁰		Hashing scenario ¹¹		Notes
	Varying MAC	Varying IP	Hashing over LAG at LSR	Hashing over ECMP paths at LSR	
2					—

¹⁰ A blank cell indicates that the MAC or IP packet header address field value used in hashing does not vary.

¹¹ A blank cell indicates that no hashing occurs for the specific scenario.

Number and types of labels egressing iLER	Packet header address fields used in hashing ¹⁰		Hashing scenario ¹¹		Notes
	Varying MAC	Varying IP	Hashing over LAG at LSR	Hashing over ECMP paths at LSR	
(LDP transport label and service label)	✓				
		✓			
	✓	✓			
3 (LDP transport label, service label, and hash label)					The last label egressing the iLER is a hash label, which has a different value from the other two (2) labels because it has different MAC and IP fields in the packet of the service traffic.
	✓		✓	✓	
		✓	✓	✓	
	✓	✓	✓	✓	
3 (LDP/RSVP transport label, BGP3107 label, and service label)					The packet egresses the LSR between the PE and ASBR with three (3) labels. Each label has the same value in every stream for traffic forwarded in a specific service. However, the values are not the same for traffic forwarded in multiple services using the same LDP LSP.
	✓				
		✓			
	✓	✓			
4 (LDP/RSVP transport label, BGP 3107 label, service label, and hash label)					The packet egresses the LSR between the PE and ASBR with three (3) labels and one (1) hash label, which is the fourth label in the packet. Hashing does not occur at the LSR between the PE and ASBR; however, if a LAG is configured on egress of the ASBR, the packets are hashed over the LAG members.
	✓				
		✓			
	✓	✓			

3.1.7 Link LDP

Hello adjacency will be brought up using link Hello packet with source IP address set to the interface borrowed IP address and a destination IP address set to 224.0.0.2.

¹⁰ A blank cell indicates that the MAC or IP packet header address field value used in hashing does not vary.

¹¹ A blank cell indicates that no hashing occurs for the specific scenario.

By default, the LDP session uses the system interface address as the LSR-ID unless explicitly configured using the command **config>router>ldp>interface-parameters>interface>local-lsr-id interface**. Using this command user is allowed to use the local interface as both the LSR-ID and the transport address for the link-level LDP session. Note that when the interface option is selected, the transport connection (TCP) for the link LDP session will also use the address of the local LDP interface as the transport address. If system is the value configured under the command **configure>router>ldp>interface-parameters>interface>transport-address**, it will be overridden.

The LSR with the highest transport address, that is, LSR-ID in this case, will bootstrap the TCP connection and LDP session.

Source and destination IP addresses of LDP packets are the transport addresses, that is, LDP LSR-IDs of systems A and B in this case.

3.1.7.1 Targeted LDP

Source and destination addresses of targeted Hello packets are the LDP LSR-IDs of systems A and B.

The user can configure the local-lsr-id option on the targeted session and change the value of the LSR-ID to either the local interface or to some other interface name, loopback or not. If the local interface is selected the IP address of the local interface will be used as the LSR-ID. In all cases, the transport address for the LDP session and the source IP address of targeted Hello message will be updated to the new LSR-ID value.

The LSR with the highest transport address, that is, LSR-ID in this case, will bootstrap the TCP connection and LDP session.

Source and destination IP addresses of LDP messages are the transport addresses, which, in this case, are the LDP LSR-IDs of systems A and B.

3.1.8 Unnumbered interface support in LDP



Note:

- This feature is supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.
- P2MP LSPs are only supported on 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T.

This feature allows LDP to establish a Hello adjacency and to resolve unicast and multicast FECs over unnumbered LDP interfaces.

This feature also extends the support of lsp-ping, p2mp-lsp-ping, and ldp-treetrace to test an LDP unicast or multicast FEC which is resolved over an unnumbered LDP interface.

3.1.8.1 Feature configuration


This feature does not introduce a new CLI command for adding an unnumbered interface into LDP. Instead, the **fec-originate** command is extended to specify the interface name, because an unnumbered interface does not have an IP address of its own. The user can, however, specify the interface name for numbered interfaces.

See the CLI section for the changes to the **fec-originate** command.

3.1.8.2 Operation of LDP over an unnumbered IP interface

Consider the setup shown in the following figure.

Figure 30: LDP adjacency and session over unnumbered interface



LSR A and LSR B have the following LDP identifiers respectively:

- <LSR Id=A> : <label space id=0>
- <LSR Id=B> : <label space id=0>

There are two P2P unnumbered interfaces between LSR A and LSR B. These interfaces are identified on each system with their unique local link identifier. In other words, the combination of {Router-ID, Local Link Identifier} uniquely identifies the interface in OSPF or IS-IS throughout the network.

A borrowed IP address is also assigned to the interface to be used as the source address of IP packets which need to be originated from the interface. The borrowed IP address defaults to the system loopback interface address, A and B respectively in this setup. The user can change the borrowed IP interface to any configured IP interface, loopback or not, by applying the following command:

config>router>if>unnumbered [<ip-int-name | ip-address>]

When the unnumbered interface is added into LDP, it will have the behavior described in the following sections.

3.1.8.2.1 Link LDP

- Hello adjacency will be brought up using link Hello packet with source IP address set to the interface borrowed IP address and a destination IP address set to 224.0.0.2.
- As a consequence of 1, Hello packets with the same source IP address should be accepted when received over parallel unnumbered interfaces from the same peer LSR-ID. The corresponding Hello adjacencies would be associated with a single LDP session.
- The transport address for the TCP connection, which is encoded in the Hello packet, will always be set to the LSR-ID of the node regardless if the user enabled the interface option under **config>router>ldp>if-params>if>ipv4>transport-address**.
- The user can configure the local-lsr-id option on the interface and change the value of the LSR-ID to either the local interface or to some other interface name, loopback or not, numbered or not. If the local interface is selected or the provided interface name corresponds to an unnumbered IP interface, the unnumbered interface borrowed IP address will be used as the LSR-ID. In all cases, the transport address for the LDP session will be updated to the new LSR-ID value but the link Hello packets will continue to use the interface borrowed IP address as the source IP address.
- The LSR with the highest transport address, that is., LSR-ID in this case, will bootstrap the TCP connection and LDP session.

- Source and destination IP addresses of LDP packets are the transport addresses, that is, LDP LSR-IDs of systems A and B in this case.

3.1.8.2.2 Targeted LDP

- Source and destination addresses of targeted Hello packet are the LDP LSR-IDs of systems A and B.
- The user can configure the local-lsr-id option on the targeted session and change the value of the LSR-ID to either the local interface or to some other interface name, loopback or not, numbered or not. If the local interface is selected or the provided interface name corresponds to an unnumbered IP interface, the unnumbered interface borrowed IP address will be used as the LSR-ID. In all cases, the transport address for the LDP session and the source IP address of targeted Hello message will be updated to the new LSR-ID value.
- The LSR with the highest transport address, That is, LSR-ID in this case, will bootstrap the TCP connection and LDP session.
- Source and destination IP addresses of LDP messages are the transport addresses, that is, LDP LSR-IDs of systems A and B in this case.

3.1.8.2.3 FEC resolution

- LDP will advertise/withdraw unnumbered interfaces using the Address/Address-Withdraw message. The borrowed IP address of the interface is used.
- A FEC can be resolved to an unnumbered interface in the same way as it is resolved to a numbered interface. The outgoing interface and next-hop are looked up in RTM cache. The next-hop consists of the router-id and link identifier of the interface at the peer LSR.
- LDP FEC ECMP next-hops over a mix of unnumbered and numbered interfaces is supported.
- All LDP FEC types are supported.
- The **fec-originate** command is supported when the next-hop is over an unnumbered interface.

All LDP features are supported except for the following:

- BFD cannot be enabled on an unnumbered LDP interface. This is a consequence of the fact that BFD is not supported on unnumbered IP interface on the system.
- As a consequence of 1, LDP FRR procedures will not be triggered via a BFD session timeout but only by physical failures and local interface down events.
- Unnumbered IP interfaces cannot be added into LDP global and peer prefix policies.

3.1.9 LDP over RSVP tunnels

LDP over RSVP-TE provides end-to-end tunnels that have two important properties, fast reroute and traffic engineering which are not available in LDP. LDP over RSVP-TE is focused at large networks (over 100 nodes in the network). Simply using end-to-end RSVP-TE tunnels will not scale. While an LER may not have that many tunnels, any transit node will potentially have thousands of LSPs, and if each transit node also has to deal with detours or bypass tunnels, this number can make the LSR overly burdened.


**Note:**

- Use of the implicit NULL MPLS label must be enabled with use of LDPoRSVP. Use the command **configure>router>rsvp>implicit-null-label** and **configure>router>ldp> implicit-null-label** to enable use of implicit NULL MPLS labels.
- Only FRR one-to-one is supported when LDPoRSVP is used. FRR facility is not supported. This is not blocked in the CLI, but operators need to ensure it when configuring the nodes.

LDP over RSVP-TE allows tunneling of user packets using an LDP LSP inside an RSVP LSP. The main application of this feature is for deployment of MPLS based services, for example, VPRN, VLL, and VPLS services, in large scale networks across multiple IGP areas without requiring full mesh of RSVP LSPs between PE routers.

The network displayed in [Figure 31: LDP over RSVP application](#) consists of two metro areas, Area 1 and 2 respectively, and a core area, Area 3. Each area makes use of TE LSPs to provide connectivity between the edge routers. To enable services between PE1 and PE2 across the three areas, LSP1, LSP2, and LSP3 are set up using RSVP-TE. There are in fact 6 LSPs required for bidirectional operation but we will refer to each bidirectional LSP with a single name, for example, LSP1. A targeted LDP (T-LDP) session is associated with each of these bidirectional LSP tunnels. That is, a T-LDP adjacency is created between PE1 and ABR1 and is associated with LSP1 at each end. The same is done for the LSP tunnel between ABR1 and ABR2, and finally between ABR2 and PE2. The loopback address of each of these routers is advertised using T-LDP. Similarly, backup bidirectional LDP over RSVP tunnels, LSP1a and LSP2a, are configured via ABR3.

Figure 31: LDP over RSVP application



This setup effectively creates an end-to-end LDP connectivity which can be used by all PEs to provision services. The RSVP LSPs are used as a transport vehicle to carry the LDP packets from one area to another. Note that only the user packets are tunneled over the RSVP LSPs. The T-LDP control messages are still sent unlabeled using the IGP shortest path.

In this application, the bidirectional RSVP LSP tunnels are not treated as IP interfaces and are not advertised back into the IGP. A PE must always rely on the IGP to look up the next hop for a service packet. LDP-over-RSVP introduces a new tunnel type, tunnel-in-tunnel, in addition to the existing LDP tunnel and RSVP tunnel types. If multiple tunnels types match the destination PE FEC lookup, LDP will prefer an LDP tunnel over an LDP-over-RSVP tunnel by default.


The design in [Figure 31: LDP over RSVP application](#) allows a service provider to build and expand each area independently without requiring a full mesh of RSVP LSPs between PEs across the three areas.

To participate in a VPRN service, PE1 and PE2 perform the autobind to LDP. The LDP label which represents the target PE loopback address is used below the RSVP LSP label. Therefore a 3 label stack is required.

To provide a VLL service, PE1 and PE2 are still required to set up a targeted LDP session directly between them. Again a 3 label stack is required, the RSVP LSP label, followed by the LDP label for the loopback address of the destination PE, and finally the pseudowire label (VC label).

This implementation supports a variation of the application in [Figure 31: LDP over RSVP application](#), in which area 1 is an LDP area. In that case, PE1 will push a two label stack while ABR1 will swap the LDP label and push the RSVP label as shown in [Figure 32: LDP over RSVP application variant](#).

Figure 32: LDP over RSVP application variant



3.1.9.1 Signaling and operation

3.1.9.1.1 LDP label distribution and FEC resolution

The user creates a targeted LDP (T-LDP) session to an ABR or the destination PE. This results in LDP hellos being sent between the two routers. These messages are sent unlabeled over the IGP path. Next, the user enables LDP tunneling on this T-LDP session and optionally specifies a list of LSP names to associate with this T-LDP session. By default, all RSVP LSPs which terminate on the T-LDP peer are candidates for LDP-over-RSVP tunnels. At this point in time, the LDP FECs resolving to RSVP LSPs are added into the Tunnel Table Manager as tunnel-in-tunnel type.

Note that if LDP is running on regular interfaces also, then the prefixes LDP learns are going to be distributed over both the T-LDP session as well as regular IGP interfaces. The policy controls which prefixes go over the T-LDP session, for example, only /32 prefixes, or a particular prefix range.

LDP-over-RSVP works with both OSPF and IS-IS. These protocols include the advertising router when adding an entry to the RTM. LDP-over-RSVP tunnels can be used as shortcuts for BGP next-hop resolution.

3.1.9.1.2 Default FEC resolution procedure

When LDP tries to resolve a prefix received over a T-LDP session, it performs a lookup in the Routing Table Manager (RTM). This lookup returns the next hop to the destination PE and the advertising router (ABR or destination PE itself). If the next-hop router advertised the same FEC over link-level LDP, LDP will prefer the LDP tunnel by default unless the user explicitly changed the default preference using the system wide prefer-tunnel-in-tunnel command. If the LDP tunnel becomes unavailable, LDP will select an LDP-over-RSVP tunnel if available.

When searching for an LDP-over-RSVP tunnel, LDP selects the advertising routers with best route. If the advertising router matches the T-LDP peer, LDP then performs a second lookup for the advertising router in the Tunnel Table Manager (TTM) which returns the user configured RSVP LSP with the best metric. If there are more than one configured LSP with the best metric, LDP selects the first available LSP.

If all user configured RSVP LSPs are down, no more action is taken. If the user did not configure any LSPs under the T-LDP session, the lookup in TTM will return the first available RSVP LSP which terminates on the advertising router with the lowest metric.

3.1.9.1.3 FEC resolution procedure when prefer-tunnel-in-tunnel is enabled

When LDP tries to resolve a prefix received over a T-LDP session, it performs a lookup in the Routing Table Manager (RTM). This lookup returns the next hop to the destination PE and the advertising router (ABR or destination PE itself).

When searching for an LDP-over-RSVP tunnel, LDP selects the advertising routers with best route. If the advertising router matches the targeted LDP peer, LDP then performs a second lookup for the advertising router in the Tunnel Table Manager (TTM) which returns the user configured RSVP LSP with the best metric. If there are more than one configured LSP with the best metric, LDP selects the first available LSP.

If all user configured RSVP LSPs are down, then an LDP tunnel will be selected if available.

If the user did not configure any LSPs under the T-LDP session, a lookup in TTM will return the first available RSVP LSP which terminates on the advertising router. If none are available, then an LDP tunnel will be selected if available.

3.1.9.2 Rerouting around failures

Every failure in the network can be protected against, except for the ingress and egress PEs. All other constructs have protection available. These constructs are LDP-over-RSVP tunnel and ABR.

3.1.9.2.1 LDP-over-RSVP tunnel protection

An RSVP LSP can deal with a failure in two ways:

- If the LSP is a loosely routed LSP, then RSVP will find a new IGP path around the failure, and traffic will follow this new path. This may involve some churn in the network if the LSP comes down and then gets re-routed. The tunnel damping feature was implemented on the LSP so that all the dependent protocols and applications do not flap unnecessarily.
- If the LSP is a CSPF-computed LSP with the fast reroute option enabled, then RSVP will switch to the detour path very quickly. From that point, a new LSP will be attempted from the head-end (global revertive). When the new LSP is in place, the traffic switches over to the new LSP with make-before-break.



Note:

Only FRR one-to-one is supported with LDP-over-RSVP with use of implicit NULL label. In other words, implicit NULL label must be enabled to use FRR one-to-one. FRR facility cannot be used. The software does not make any checks to enforce these restrictions. Operators must ensure this by network design and configuration.

3.1.9.2.2 ABR protection

If an ABR fails, then routing around the ABR requires that a new next-hop LDP-over-RSVP tunnel be found to a backup ABR. If an ABR fails, then the T-LDP adjacency fails. Eventually, the backup ABR becomes the new next hop (after SPF converges), and LDP learns of the new next-hop and can reprogram the new path.

3.1.10 T-LDP session tracking using BFD

The user enables BFD tracking of a T-LDP session by using the **config>router>ldp>targeted-session>bfd-enable** command.

When this command is executed, LDP registers the address of the T-LDP session peer with BFD for tracking purposes. In other words, when the BFD session goes down, the T-LDP session is also brought down. However, the BFD session going up does not affect the state of the T-LDP session as T-LDP has to establish correct Hello adjacency and then a TCP connection to the peer which then allows the T-LDP session to come up.

The source and destination addresses of the BFD session depends on whether the T-LDP peer is directly reachable over a local interface or is more than one hop away.

When the peer is on the local subnet, the BFD session used will be the one associated with the local interface on the direct link to the peer. In that case, the source address and destination address in the BFD packets will be that of the local end and the far-end of that interface respectively. If multiple interfaces exist to the peer because of parallel links, then the BFD session must be associated with the interface which is currently used by the common LDP session shared by both the T-LDP and link-level LDP sessions.

The parameters used for the BFD session, **transmit-interval**, **receive-interval**, **multiplier**, and **echo-receive** are also configured under the local interfaces using the **config>router>interface>bfd** command.

Note that the local interface BFD session is used regardless if the LDP session, and underlying TCP connection, were bootstrapped by the link-level LDP Hello adjacency or the T-LDP hello adjacency. Furthermore, if the BFD session goes down it will bring down the state of both the T-LDP session and the link-level LDP session sharing the same LDP session.

When the peer is several hops away, the BFD session used will be the one associated with the loopback interface corresponding to LSR-ID of the T-LDP session. The LSR-ID is used to establish the Hello adjacency with the peer. By default the LSR-ID matches the system interface address but the user can change it to any other loopback interface address [ldp-instances]. In that case, the source address and destination address in the BFD packets will match the local end LSR-ID and the far-end address specified for the peer respectively. The parameters used for the BFD session are also those configured under the loopback interface corresponding to the LSR-ID using the **bfd** command in the **config>router>interface** context.

Because the BFD session used to track the same T-LDP peer may move from a link interface to a loopback interface depending on route reachability, it is important that the user configures the BFD session parameters consistently on both interfaces.

The link interface BFD session is sourced and maintained on the IOM while the loopback interface BFD session is sourced and maintained on the CPM. As a result, the system level BFD resource count reflects the worst case where each T-LDP session is using two BFD sessions.

3.1.10.1 LDP Downstream-on-Demand (DoD)

The user enables the use by an LDP session of the Downstream-on-Demand (DoD) label distribution using the command **config>router>ldp>peer-parameters>peer> dod-label-distribution**.

When this option is enabled, LDP will set the A-bit in the Label Initialization message, when the LDP session to the peer is established. When both peers set the A-bit, both uses the DoD label distribution method over the LDP session [rfc5036].

This feature can only be enabled on a link level LDP session and applies to prefix labels only, and not service labels.

3.1.10.1.1 Single-hop LDP DoD procedures

As soon as the link LDP session comes up, the 7210 SAS sends a label request to the DoD peer for the FEC prefix corresponding to the peer's LSR-id. The DoD peer LSR-id is found in the basic Hello discovery messages the peer used to establish the Hello adjacency with the 7210.

Similarly, if the 7210 SAS and the directly attached DoD peer enter into the extended discovery and establish a targeted LDP session, the 7210 SAS immediately sends a label request for the FEC prefix corresponding to the peer's LSR-id found in the extended discovery messages.

However, the 7210 SAS node does not advertise any <FEC, label> bindings, including the FEC of its own LSR-id, unless the DoD peer requested it through a Label Request Message.

When the DoD peer sends a label request for any FEC prefix, the 7210 SAS replies with a <FEC, label> binding for that prefix if the FEC was already activated on the 7210 SAS. If not, the 7210 SAS replies with a notification message containing the status code of "no route". The 7210 SAS does not attempt in the latter case to send a label request to the next-hop for the FEC prefix when the LDP session to this next-hop uses the DoD label distribution mode. Thus, the reference to single-hop LDP DoD procedures.

The single-hop LDP DoD procedures makes sure the 7210 SAS has a label for the LDP DoD peer, whenever it is needed.

The 7210 SAS needs a label of directly attached DoD peer in the following cases:

- A BGP labeled route for the peer's prefix from RTM to its BGP neighbors through iBGP.
- When it receives a label request message from a directly attached DoD peer for the prefix of another directly attached DoD peer. In this case the DoD peers are trying to establish a SDP among themselves.
- Trying to establish a SDP to a directly attached LDP DoD peer.

The 7210 SAS also supports sending and receiving the Label Abort Request Message as described. This message is used to abort an outstanding request for a label in case no response was received from the peer within a finite amount of time.

3.1.11 LDP over RSVP and ECMP

7210 SAS devices does not support ECMP for LDP over RSVP LSPs.

3.1.12 LDP Fast-Reroute for IS-IS and OSPF prefixes

LDP Fast Re-Route (FRR) is a feature which allows the user to provide local protection for an LDP FEC by precomputing and downloading to IOM both a primary and a backup NHLFE for this FEC.

The primary NHLFE corresponds to the label of the FEC received from the primary next-hop as per standard LDP resolution of the FEC prefix in RTM. The backup NHLFE corresponds to the label received for the same FEC from a Loop-Free Alternate (LFA) next-hop.

The LFA next-hop precomputation by IGP is described in RFC 5286 – “Basic Specification for IP Fast Reroute: Loop-Free Alternates”. LDP FRR relies on using the label-FEC binding received from the LFA next-hop to forward traffic for a specified prefix as soon as the primary next-hop is not available. This means that a node resumes forwarding LDP packets to a destination prefix without waiting for the routing convergence. The label-FEC binding is received from the loop-free alternate next-hop ahead of time and is stored in the Label Information Base because LDP on the router operates in the liberal retention mode.

This feature requires that IGP performs the Shortest Path First (SPF) computation of an LFA next-hop, in addition to the primary next-hop, for all prefixes used by LDP to resolve FECs. IGP also populates both routes in the Routing Table Manager (RTM).

3.1.12.1 LDP FRR configuration

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

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

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

The above 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.

Next the user enables the use by LDP of the LFA next-hop by configuring the following option:

```
config>router>ldp>fast-reroute
```

When this command is enabled, LDP will use both the primary next-hop and LFA next-hop, when available, for resolving the next-hop of an LDP FEC against the corresponding prefix in the RTM. This will result in LDP programming a primary NHLFE and a backup NHLFE into the IOMXCM for each next-hop of a FEC prefix for the purpose of forwarding packets over the LDP FEC.

Note that because LDP can detect the loss of a neighbor/next-hop independently, it is possible that it switches to the LFA next-hop while IGP is still using the primary next-hop. To avoid this situation, it is recommended to enable IGP-LDP synchronization on the LDP interface:

```
config>router>interface>ldp-sync-timer seconds
```

3.1.12.1.1 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
```


Note that if IGP shortcut are also enabled in LFA SPF, LSPs with destination address in that IS-IS level or OSPF area are also not included in the LFA SPF calculation.

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 above OSPF command can only be executed under the area in which the specified interface is primary and when enabled, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to apply it to an area where the interface is secondary, the command will fail.

3.1.12.2 LDP FRR procedures

The LDP FEC resolution when LDP FRR is not enabled operates as follows. When LDP receives a FEC, label binding for a prefix, it will resolve it by checking if the exact prefix, or a longest match prefix when the **aggregate-prefix-match option** is enabled in LDP, exists in the routing table and is resolved against a next-hop which is an address belonging to the LDP peer which advertised the binding, as identified by its LSR-id. When the next-hop is no longer available, LDP deactivates the FEC and deprograms the NHLFE in the datapath. LDP will also immediately withdraw the labels it advertised for this FEC and deletes the ILM in the datapath unless the user configured the **label-withdrawal-delay** option to delay this operation. Traffic that is received while the ILM is still in the datapath is dropped. When routing computes and populates the routing table with a new next-hop for the prefix, LDP resolves again the FEC and programs the datapath accordingly.

When LDP FRR is enabled and an LFA backup next-hop exists for the FEC prefix in RTM, or for the longest prefix the FEC prefix matches to when **aggregate-prefix-match** option is enabled in LDP, LDP will resolve the FEC as above but will program the datapath with both a primary NHLFE and a backup NHLFE for each next-hop of the FEC.

In order perform a switchover to the backup NHLFE in the fast path, LDP follows the uniform FRR failover procedures which are also supported with RSVP FRR.

When any of the following events occurs, LDP instructs in the fast path the IOM to enable the backup NHLFE for each FEC next-hop impacted by this event. The IOM do that by simply flipping a single state bit associated with the failed interface or neighbor/next-hop:

1. An LDP interface goes operationally down, or is admin shutdown. In this case, LDP sends a neighbor/next-hop down message to the IOM for each LDP peer it has adjacency with over this interface.
2. An LDP session to a peer went down as the result of the Hello or Keep-Alive timer expiring over a specific interface. In this case, LDP sends a neighbor/next-hop down message to the IOM for this LDP peer only.
3. The TCP connection used by a link LDP session to a peer went down, due say to next-hop tracking of the LDP transport address in RTM, which brings down the LDP session. In this case, LDP sends a neighbor/next-hop down message to the IOM for this LDP peer only.
4. A BFD session, enabled on a T-LDP session to a peer, times-out and as a result the link LDP session to the same peer and which uses the same TCP connection as the T-LDP session goes also down. In this case, LDP sends a neighbor/next-hop down message to the IOM for this LDP peer only.

5. A BFD session enabled on the LDP interface to a directly connected peer, times-out and brings down the link LDP session to this peer. In this case, LDP sends a neighbor/next-hop down message to the IOM for this LDP peer only. BFD support on LDP interfaces is a new feature introduced for faster tracking of link LDP peers.

The tunnel-down-dump-time option or the label-withdrawal-delay option, when enabled, does not cause the corresponding timer to be activated for a FEC as long as a backup NHLFE is still available.

3.1.12.2.1 Link LDP Hello adjacency tracking with BFD

LDP can only track an LDP peer with which it established a link LDP session with using the Hello and Keep-Alive timers. If an IGP protocol registered with BFD on an IP interface to track a neighbor, and the BFD session times out, the next-hop for prefixes advertised by the neighbor are no longer resolved. This however does not bring down the link LDP session to the peer because the LDP peer is not directly tracked by BFD. More importantly the LSR-id of the LDP peer may not coincide with the neighbor's router-id IGP is tracking by way of BFD.

To properly track the link LDP peer, LDP needs to track the Hello adjacency to its peer by registering with BFD. This way, the peer next-hop is tracked.

The user enables Hello adjacency tracking with BFD by enabling BFD on an LDP interface:

```
config>router>ldp>interface-parameters>interface>enable-bfd
```

The parameters used for the BFD session, that is, transmit-interval, recethat is interval, and multiplier, are those configured under the IP interface in existing implementation:

```
config>router>interface>bfd
```

When multiple links exist to the same LDP peer, a Hello adjacency is established over each link but only a single LDP session will exist to the peer and will use a TCP connection over one of the link interfaces. Also, a separate BFD session should be enabled on each LDP interface. If a BFD session times out on a specific link, LDP will immediately bring down the Hello adjacency on that link. In addition, if there are FECs which have their primary NHLFE over this link, LDP triggers the LDP FRR procedures by sending to IOM the neighbor/next-hop down message. This will result in moving the traffic of the impacted FECs to an LFA next-hop on a different link to the same LDP peer or to an LFA backup next-hop on a different LDP peer depending on the lowest backup cost path selected by the IGP SPF.

As soon as the last Hello adjacency goes down because of BFD timing out, the LDP session goes down and the LDP FRR procedures will be triggered. This will result in moving the traffic to an LFA backup next-hop on a different LDP peer.

3.1.12.2.2 ECMP considerations

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

Also note that when the system ECMP value is set to **ecmp=1** or to **no ecmp**, which translates to the same and is the default value, SPF will be able to use the overflow ECMP links as LFA next hops in these two cases.

3.1.13 LDP P2MP support

This section describes support for LDP P2MP.

3.1.13.1 LDP P2MP configuration



Note:

- This feature is supported on all 7210 SAS platforms as described in this document, except the 7210 SAS-Sx 1/10GE and 7210 SAS-Sx 10/100GE, and platforms operating in access-uplink mode.
- P2MP LSPs signaled using RSVP or mLDP is only supported on 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

A node running LDP also supports P2MP LSP setup using LDP. By default, it would advertise the capability to a peer node using P2MP capability TLV in LDP initialization message.

This configuration option per interface is provided to restrict/allow the use of interface in LDP multicast traffic forwarding toward a downstream node. The interface configuration option does not restrict/allow exchange of P2MP FEC by way of established session to the peer on an interface, but it would only restrict/allow use of next-hops over the interface. By default, the LDP-P2MP capability is disabled on interface.

3.1.13.2 LDP P2MP protocol

Only a single generic identifier range is defined for signaling multipoint tree for all client applications. Implementation on 7210 SAS reserves the range (1 to 8292) of generic LSP P2MP-ID on root node for static P2MP LSP.

3.1.13.3 Configuration guidelines for P2MP LSPs


- Before using P2MP LSPs with NG-MVPN, resources must be allocated from the sf-ingress-internal-tcam resource pool using the **configure>system>global-res-profile>sf-ingress-internal-tcam>mpls-p2mp** command. In addition, if the 7210 SAS-R6 is deployed as a bud router, the **configure>system>loopback-no-svc-port p2mpbud p2mpbud-port-id** command must be used to configure one of the front-panel ports as a loopback port.
- Ingress FC classification is available for packets received on a P2MP LSP on a network port IP interface that needs to be replicated to IP receivers. Ingress FC classification allows users to prioritize multicast traffic to IP receivers in the service. Also available is the capability to mark the packet with IP DSCP values while sending the multicast stream out of the IP interface. To enable ingress FC classification, use the **loopback-no-svc-port [p2mpbud p2mpbud-port-id [classification]]** command. Before using the command, users must ensure that sufficient resources are available in the network port ingress CAM resource pool and MPLS EXP ingress profile map resource pool. The **tools>dump>system-resources** command can be used to check resource availability.

3.1.14 IS-IS and OSPF 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 33: Topology with primary and LFA routes




al_0215

The primary route is by way of R3. The LFA route by way of 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. R2 must be loop-free with respect to source node R1.

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. Thus 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. In other words it is required that R2 becomes loop-free with respect to both the source node R1 and the protected node R3.

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

Figure 34: Example topology with broadcast interfaces



al_0216

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's Pseudo-Node (PN). However, since R2 has also a link to that PN, its cost to reach R5 by way of the PN or router R4 are the same. Thus 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 rules for this criterion as provided in *RFC 5286*:

- **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 specified primary next-hop, it follows the following selection algorithm:

1. It will pick the node-protect type in favor of the link-protect type.
2. If there is more than one LFA next-hop within the selected type, then it will pick one based on the least cost.
3. If more than one LFA next-hop with the same cost results from Step B, then SPF will select the first one. This is not a deterministic selection and will vary following each SPF calculation.

For the case of a broadcast interface, a node-protect LFA is not necessarily a link protect LFA if the path to the LFA next-hop goes over the same PN as the primary next-hop. Similarly, a link protect LFA may not guarantee link protection if it goes over the same PN as the primary next-hop.

The selection algorithm when SPF finds multiple LFA next-hops for a specified 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 go over the PN used by the primary next-hop.
2. If there is more than one LFA next-hop in the first set, it will pick the node-protect type in favor of the link-protect type.
3. If there is more than one LFA next-hop within the selected type, then it will pick one based on the least cost.
4. If more than one LFA next-hop with equal cost results from Step C, SPF will select the first one from the remaining set. This is not a deterministic selection and will vary following each SPF calculation.
5. If no LFA next-hop results from Step D, SPF will rerun Steps B-D using the second set.



Note:

This algorithm is more flexible than strictly applying Rule 3 above; 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; 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 specified prefix are programmed into RTM.

3.1.14.1 Loop-Free Alternate calculation for inter-area/inter-level prefixes

When SPF resolves OSPF inter-area prefixes or IS-IS inter-level prefixes, it will compute an LFA backup next-hop to the same exit area/border router as used by the primary next-hop.

3.1.14.2 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. See more details in the Loop-Free Alternate Shortest Path First (LFA SPF) Policies section in the *7210 SAS-Mxp, R6, R12, S, Sx, T Routing Protocols Guide*.

3.1.15 Multi-area and multi-instance extensions to LDP

To extend LDP across multiple areas of an IGP instance or across multiple IGP instances, the current standard LDP implementation based on RFC 3036 requires that all the /32 prefixes of PEs be leaked between the areas or instances. This is because an exact match of the prefix in the routing table has to install the prefix binding in the LDP Forwarding Information Base (FIB).

Multi-area and multi-instance extensions to LDP provide an optional behavior by which LDP installs a prefix binding in the LDP FIB by simply performing a longest prefix match with an aggregate prefix in the routing table (RIB). The ABR is configured to summarize the /32 prefixes of PE routers. This method is compliant to RFC 5283- LDP Extension for Inter-Area Label Switched Paths (LSPs).

3.2 LDP IPv6 control and data planes

**Note:**

LDP IPv6 is supported only on the 7210 SAS-Mxp.


7210 SAS extends the LDP control and data planes to support the LDP IPv6 adjacency and session using 128-bit LSR-ID.

The implementation provides concurrent support of independent LDP IPv4 (32-bit LSR-ID) and IPv6 (128-bit LSR-ID) adjacencies and sessions between peer LSRs and over the same or different set of interfaces.

3.2.1 LDP Operation in an IPv6 Network

LDP IPv6 can be enabled on the 7210 SAS interface. The following figure shows the LDP adjacency and session over an IPv6 interface.

Figure 35: LDP adjacency and session over an IPv6 interface



In the preceding example, LSR-A and LSR-B have the following IPv6 LDP identifiers respectively:

- <LSR Id=A/128> : <label space id=0>
- <LSR Id=B/128> : <label space id=0>

By default, A/128 and B/128 use the system interface IPv6 address.

**Note:**

Although the LDP control plane can operate using only the IPv6 system address, for optimal operation, the user must configure the IPv4-formatted router ID for OSPF, IS-IS, and BGP.

3.2.2 Link LDP

The 7210 SAS LDP IPv6 implementation uses a 128-bit LSR-ID as defined in *draft-pdutta-mpls-ldp-v2-00*. See [LDP process overview](#) for more information about interoperability of this implementation with 32-bit LSR-ID, as defined in RFC 7552.

Hello adjacency is brought up using a link Hello packet with the source IP address set to the interface link-local unicast address and a destination IP address set to the link-local multicast address FF02:0:0:0:0:0:2.

The transport address for the TCP connection, which is encoded in the Hello packet, is set to the LSR-ID of the LSR by default. It is set to the interface IPv6 address if the user enables the **interface** option under one of the following contexts:

- **config>router>ldp>if-params>ipv6>transport-address**
- **config>router>ldp>if-params>if>ipv6>transport-address**

The interface global unicast address, that is, the primary IPv6 unicast address of the interface, is used.

The user can configure the **local-lsr-id** option on the interface and change the value of the LSR-ID to either the local interface or to another interface name, whether a loopback interface or any other non-loopback interface. The global unicast IPv6 address corresponding to the primary IPv6 address of the interface is used as the LSR-ID. If the user invokes an interface that does not have a global unicast IPv6 address in the configuration of the transport address or the configuration of the **local-lsr-id** option, the session does not come up and an error message is displayed.

The LSR with the highest transport address bootstraps the IPv6 TCP connection and IPv6 LDP session.

Source and destination addresses of LDP and TCP session packets are the IPv6 transport addresses.

3.2.3 Targeted LDP

Source and destination addresses of the targeted Hello packet are the LDP IPv6 LSR IDs of systems A and B, as shown in [Figure 35: LDP adjacency and session over an IPv6 interface](#).

The user can configure the **local-lsr-id** option on the targeted session and change the value of the LSR ID to either the local interface or to some other interface name, whether a loopback interface or any other non-loopback interface. The global unicast IPv6 address corresponding to the primary IPv6 address of the interface is used as the LSR ID. If the user invokes an interface that does not have a global unicast IPv6 address in the configuration of the transport address or the configuration of the **local-lsr-id** option, the session does not come up and an error message is displayed. In all cases, the transport address for the LDP session and the source IP address of the targeted Hello message is updated to the new LSR ID value.

The LSR with the highest transport address (in this case, the LSR ID) bootstraps the IPv6 TCP connection and IPv6 LDP session.

Source and destination IP addresses of LDP and TCP session packets are the IPv6 transport addresses (in this case, LDP LSR IDs of systems A and B).

3.2.4 FEC resolution

LDP advertises and withdraws all interface IPv6 addresses using the Address and Address-Withdraw message. Both the link-local unicast address and the configured global unicast addresses of an interface are advertised.

All LDP FEC types can be exchanged over an LDP IPv6 LDP session, similar to an LDP IPv4 session.

The LSR does not advertise a FEC for a link-local address and, if received, the LSR does not resolve it.

An IPv4 or IPv6 prefix FEC can be resolved to an LDP IPv6 interface in the same way as it is resolved to an LDP IPv4 interface. The outgoing interface and next hop are looked up in the RTM cache. The next hop can be the link-local unicast address of the other side of the link or a global unicast address. The FEC is resolved to the LDP IPv6 interface of the downstream LDP IPv6 LSR that advertised the IPv4 or IPv6 address of the next hop.

An mLDP P2MP IPv4 FEC with an IPv4 root LSR PE, and carrying one or more IPv4 multicast prefixes, can be resolved to an upstream LDP IPv6 LSR by checking if the LSR advertised the next hop for the IPv4 root PE. The upstream LDP IPv6 LSR then resolves the IPv4 P2MP FEC to one of the LDP IPv6 links to this LSR.



Note:

- The 7210 SAS does not support IPv6 multicast and the use of P2MP LSPs for IPv6 multicast.
- Manually configured mLDP P2MP LSPs, NG-mVPN, and dynamic mLDP cannot operate in an IPv6-only network.

A PW FEC can be resolved to a targeted LDP IPv6 adjacency with an LDP IPv6 LSR if there is a context for the FEC with local spoke-SDP configuration.

3.2.5 Resources required to trap LDP control packets to the CPU

On the 7210 SAS, by default, resources required to send LDP IPv6 packets for control plane processing are not available. The resources are shared with other control packets, and the total number of resources is configured using the **config>system>resource-profile>ingress-internal-tcam>ip-mpls-protocols** command. The user must allocate the appropriate number of resources using the **ip-mpls-protocols** command to ensure that LDP IPv6 can be configured for use.

The LDP IPv6 IFP entry is installed when an interface is enabled for LDP IPv6. The user can configure the **config>router>ldp>interface-parameters>interface>ipv6** command to trigger the installation of the IFP entry.

The IFP entry is installed for the first LDP IPv6-enabled interface. It is uninstalled when the last interface is LDP IPv6-disabled.



Note:

- The LDP IPv6 IFP entry is uninstalled only when the user configures the **configure>router>ldp>interface-parameters>interface>no ipv6** command for the last interface in the **ldp** context. Executing the **configure>router>ldp>interface-parameters>interface>ipv6>shutdown** command on the last interface does not uninstall the IFP entry.

- The user can issue the **config>router>no ldp** command only after issuing the **config>router>ldp>shutdown** command.

3.2.6 LDP session capabilities

LDP supports advertisement of all FEC types over an LDP IPv4 or an LDP IPv6 session. These FEC types are: IPv4 prefix FEC, IPv6 prefix FEC, IPv4 P2MP FEC, PW FEC 128, and PW FEC 129.

In addition, LDP supports signaling the enabling or disabling of the advertisement of the following subset of FEC types, both during the LDP IPv4 or IPv6 session initialization phase, and subsequently when the session is already up.

- **IPv4 prefix FEC**

This is performed using the State Advertisement Control (SAC) capability TLV, as described in RFC 7473. The SAC capability TLV includes the IPv4 SAC element having the Disable-bit (D-bit) set or reset to disable or enable this FEC type, respectively. The LSR can send this TLV in the LDP Initialization message and subsequently in an LDP Capability message.

- **IPv6 prefix FEC**

This is performed using the SAC capability TLV, as described in RFC 7473. The SAC capability TLV includes the IPv6 SAC element having the D-bit set or reset to disable or enable this FEC type, respectively. The LSR can send this TLV in the LDP Initialization message and subsequently in a LDP Capability message to update the state of this FEC type.

- **P2MP FEC**

This is performed using the P2MP capability TLV, as described in RFC 6388. The P2MP capability TLV has the State-bit (S-bit) with a value of set or reset to enable or disable this FEC type, respectively. Unlike the IPv4 SAC and IPv6 SAC capabilities, the P2MP capability does not distinguish between the IPv4 and IPv6 P2MP FEC. The LSR can send this TLV in the LDP Initialization message and, subsequently, in a LDP Capability message to update the state of this FEC type.

During LDP session initialization, each LSR indicates to its peers the FEC type it supports by including the capability TLV for it in the LDP Initialization message. The 7210 SAS implementation enables the preceding FEC types by default and consequently sends the corresponding capability TLVs in the LDP initialization message. If one or both peers advertise the disabling of a capability in the LDP Initialization message, no FECs of the corresponding FEC type are exchanged between the two peers for the lifetime of the LDP session unless a Capability message is sent subsequently to explicitly enable it. The same behavior applies if no capability TLV for a FEC type is advertised in the LDP initialization message, except for the IPv4 prefix FEC, which is assumed to be supported by all implementations by default.

The Dynamic Capability, as defined in RFC 5561, allows all the preceding FEC types to update the enabled or disabled state after the LDP session initialization phase. An LSR informs its peer that it supports the Dynamic Capability by including the Dynamic Capability Announcement TLV in the LDP Initialization message. If both LSRs advertise this capability, the user is allowed to enable or disable any of the preceding FEC types while the session is up and the change takes effect immediately. The LSR then sends a SAC Capability message with the IPv4 or IPv6 SAC element having the D-bit set or reset, or the P2MP capability TLV in a Capability message with the S-bit set or reset. Each LSR then takes the consequent action of withdrawing or advertising the FECs of that type to the peer LSR. If one or both LSRs did not advertise the Dynamic Capability Announcement TLV in the LDP Initialization message, any change to the enabled or disabled FEC types only takes effect the next time the LDP session is restarted.

The user can enable or disable a specific FEC type for an LDP session to a peer using the following CLI commands:

- **config>router>ldp>session-params>peer>fec-type-capability p2mp**
- **config>router>ldp>session-params>peer>fec-type-capability prefix-ipv4**
- **config>router>ldp>session-params>peer>fec-type-capability prefix-ipv6**

3.2.7 LDP adjacency capabilities

Adjacency-level FEC-type capability advertisement is defined in *draft-pdutta-mpls-ldp-adj-capability*. By default, all FEC types supported by the LSR are advertised in the LDP IPv4 or IPv6 session initialization; see [LDP session capabilities](#) for more information. If a specific FEC type is enabled at the session level, it can be disabled over a specific LDP interface at the IPv4 or IPv6 adjacency level for all IPv4 or IPv6 peers over that interface. If a specific FEC type is disabled at the session level, FECs are not advertised and enabling that FEC type at the adjacency level does not have any effect. The LDP adjacency capability can be configured on link Hello adjacency only and does not apply to targeted Hello adjacency.

The LDP adjacency capability TLV is advertised in the Hello message with the D-bit set or reset to disable or enable the resolution of this FEC type over the link of the Hello adjacency. It is used to restrict which FECs can be resolved over a specific interface to a peer. This provides the ability to dedicate links and data path resources to specific FEC types. An mLDP P2MP FEC can exclude specific links to a downstream LSR from being used to resolve this type of FEC.

Like the LDP session-level FEC-type capability, the adjacency FEC-type capability is negotiated for both directions of the adjacency. If one or both peers advertise disabling a capability in the LDP Hello message, no FECs of the corresponding FEC type are resolved by either peer over the link of this adjacency for the lifetime of the LDP Hello adjacency, unless one or both peers send the LDP adjacency capability TLV subsequently to explicitly enable it.

The user can enable or disable a specific FEC type for an LDP interface to a peer using the following CLI commands:

- **config>router>ldp>if-params>if>ipv4/ipv6>fec-type-capability p2mp-ipv4**
- **config>router>ldp>if-params>if>ipv4/ipv6>fec-type-capability prefix-ipv4**
- **config>router>ldp>if-params>if> ipv4/ipv6>fec-type-capability prefix-ipv6**

These commands, when applied for the P2MP FEC, deprecate the existing command **multicast-traffic {enable | disable}** under the interface. Unlike the session-level capability, these commands can disable multicast FEC for IPv4 and IPv6 separately.

The encoding of the adjacency capability TLV uses a private Vendor TLV. It is used only in a Hello message to negotiate a set of capabilities for a specific LDP IPv4 or IPv6 Hello adjacency.

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|1|0| ADJ_CAPABILITY_TLV          | Length          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               VENDOR_OUI          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|S| Reserved                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Adjacency capability elements
+
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The value of the U-bit for the TLV is set to 1 so that a receiver must silently ignore if the TLV is deemed unknown.

The value of the F-bit is 0. After being advertised, this capability cannot be withdrawn; therefore, the S-bit is set to 1 in a Hello message.

Adjacency capability elements are encoded as follows:

```

0 1 2 3 4 5 6 7
+---+---+---+---+
|D|  CapFlag    |
+---+---+---+---+

```

- D bit: Controls the capability state.
- 1 : Disable capability
- 0 : Enable capability
- CapFlag: The adjacency capability
- 1 : Prefix IPv4 forwarding
- 2 : Prefix IPv6 forwarding
- 3 : P2MP IPv4 forwarding
- 5 : MP2MP IPv4 forwarding

Each CapFlag appears no more than once in the TLV. If duplicates are found, the D-bit of the first element is used. For forward compatibility, if the CapFlag is unknown, the receiver must silently discard the element and continue processing the rest of the TLV.

3.2.8 Address and FEC distribution

After an LDP LSR initializes the LDP session to the peer LSR and the session comes up, local IPv4 and IPv6 interface addresses are exchanged using the Address and Address Withdraw messages. Similarly, FECs are exchanged using Label Mapping messages.

By default, IPv6 address distribution is determined by whether the Dual-stack capability TLV, which is defined in RFC 7552, is present in the Hello message from the peer. This coupling is introduced to prevent interoperability issues with existing third-party LDP IPv4 implementations.


The following is the detailed behavior for the processing of the Dual-stack capability TLV in conjunction with the IPv6 SAC TLV in the Hello message:

- If the peer has sent the dual-stack capability TLV in the Hello message, IPv6 local addresses are sent to the peer. The user can configure a new address export policy to further restrict the local IPv6 interface addresses sent to the peer. If the peer explicitly enabled the LDP IPv6 FEC type by including the IPv6 SAC TLV with the D-bit set to 0 in the initialization message, IPv6 FECs are sent to the peer. The FEC prefix export policies can be used to restrict the LDP IPv6 FEC that can be sent to the peer.
- If the peer has sent the dual-stack capability TLV in the Hello message, but explicitly disabled the LDP IPv6 FEC type by including the IPv6 SAC TLV with the D-bit set to 1 in the initialization message, IPv6 FECs are not sent, but IPv6 local addresses are sent to the peer. A CLI is provided to allow the configuration of an address export policy to further restrict the local IPv6 interface addresses that can be sent to the peer. The FEC prefix export policy has no effect because the peer has explicitly requested disabling the IPv6 FEC type advertisement.

- If the peer has not sent the dual-stack capability TLV in the Hello message, no IPv6 addresses or IPv6 FECs are sent to that peer, regardless of whether the IPv6 SAC TLV is present in the initialization message. This case is added to prevent interoperability issues with existing third-party LDP IPv4 implementations. The user can override this by explicitly configuring an address export policy and a FEC export policy to select the addresses and FECs to send to the peer.


The preceding behavior applies to LDP IPv4 and IPv6 addresses and FECs. The procedure is summarized in the flowcharts shown in the following figures.

Figure 36: LDP IPv6 address and FEC distribution procedure



al_0625

Figure 37: LDP IPv6 address and FEC distribution procedure



3.2.9 Controlling IPv6 FEC distribution during an upgrade to 7210 SAS SR OS supporting LDP IPv6

About this task

A FEC for each IPv4 and IPv6 system interface address is advertised and resolved automatically by the LDP peers when the LDP session comes up, regardless of whether the session is IPv4 or IPv6.

To avoid the automatic advertisement and resolution of an IPv6 system FEC when the LDP session is IPv4, perform this procedure before and after the upgrade to the 7210 SAS SR OS version that introduces support of LDP IPv6.

Procedure

- Step 1.** Before the upgrade, implement a global prefix policy that rejects prefix `::0/0 longer` to prevent IPv6 FECs from being installed after the upgrade.
- Step 2.** In the cold upgrade case:
- If new IPv4 sessions are created on the node, the per-peer FEC capabilities must be configured to filter out IPv6 FECs.
 - On older, pre-existing IPv4 sessions, the per-peer FEC-capabilities must be configured to filter out IPv6 FECs.
- Step 3.** When all LDP IPv4 sessions have dynamic capabilities enabled, with per-peer FEC-capabilities for IPv6 FECs disabled, the global import policy can be removed.


3.2.10 Handling of duplicate link-local IPv6 addresses in FEC resolution

Link-local IPv6 addresses are scoped to a link and, consequently, duplicate addresses can be used on different links to the same or different peer LSRs. When the duplicate addresses exist on the same LAN, routing detects them and blocks one of them. In all other cases, duplicate links are valid because they are scoped to the local link.

In this section, LLn refers to Link-Local address (n).

The following figure shows FEC resolution in a LAN.


Figure 38: FEC resolution in LAN



LSR B resolves an mLDP FEC with the root node being Root LSR. The route lookup indicates that the best route to loopback of Root LSR is {interface if-B and next-hop LL1}.

However, LDP finds that both LSR A and LSR C have advertised address LL1 and that there are Hello adjacencies (IPv4 or IPv6) to both A and C. In this case, a change is made so that an LSR only advertises link-local IPv6 addresses to a peer for the links over which it established a Hello adjacency to that peer. In this case, LSR C advertises LL1 to LSR E, but not to LSRs A, B, and D. This behavior applies with P2P and broadcast interfaces.

The preceding solution also applies if ambiguity exists with prefix FEC (unicast FEC).

Example: FEC resolution over P2P links

```
| - (LL3) - - - - - |
```

LSR B resolves an mLDP FEC with the root node being Root LSR. The route lookup indicates that the best route to loopback of Root LSR is {interface if-B and next-hop LL1}. The following describes the FEC resolution use cases:

- **Case 1: LDP is enabled on all links**

This case has no ambiguity. LDP only selects LSR A because the address LL1 from LSR C is discovered over a different interface. This case also applies to prefix FEC (unicast FEC), and consequently there is no ambiguity in the resolution.

- **Case 2: LDP is disabled on link A-B with next-hop LL1**

LSR B can still select one of the other two interfaces to upstream LSR A, as long as LSR A has advertised the LL1 address in the LDP session.

3.2.11 IGP and static route synchronization with LDP

The IGP-LDP synchronization and the static route to LDP synchronization features are modified to operate on a dual-stack IPv4 or IPv6 LDP interface as follows:

- If the router interface goes down or both LDP IPv4 and LDP IPv6 sessions go down, IGP sets the interface metric to the maximum value and all static routes with the **ldp-sync** option enabled and resolved on this interface are deactivated.
- If the router interface is up and only one of the LDP IPv4 or LDP IPv6 interfaces goes down, no action is taken.
- When the router interface comes up from a down state, and one of either the LDP IPv4 or LDP IPv6 sessions comes up, IGP starts the synchronization timer at the expiry of which the interface metric is restored to its configured value. All static routes with the **ldp-sync** option enabled are also activated at the expiry of the timer.

Because of the preceding behavior, Nokia recommends that the user should configure the synchronization timer to a value that allows enough time for both the LDP IPv4 and LDP IPv6 sessions to come up.

3.2.12 BFD operation

The operation of BFD over a LDP interface tracks the next hop of prefix IPv4 and prefix IPv6, in addition to tracking of the LDP peer address of the Hello adjacency over that link. This tracking is required as LDP can now resolve both IPv4 and IPv6 prefix FECs over a single IPv4 or IPv6 LDP session and, therefore, the next hop of a prefix does not necessarily match the LDP peer source address of the Hello adjacency. The failure of either or both of the BFD session tracking the FEC next hop and the one tracking the Hello adjacency cause the LFA backup NHLFE for the FEC to be activated, or the FEC to be re-resolved if there is no FRR backup.

The following CLI command provides the option to track only with an IPv4 BFD session, only with an IPv6 BFD session, or both:

```
config>router>ldp>if-params>if>bfd-enable [ipv4] [ipv6]
```

This command provides the flexibility required in case the user does not need to track both the Hello adjacency and next hops of FECs. For example, if the user configures **bfd-enable ipv6** only to save on the number of BFD sessions, LDP tracks the IPv6 Hello adjacency and the next hops of IPv6 prefix FECs. LDP does not track next hops of IPv4 prefix FECs resolved over the same LDP IPv6 adjacency. If the IPv4

data plane encounters errors but the IPv6 Hello adjacency is unaffected and remains up, traffic for the IPv4 prefix FECs resolved over that IPv6 adjacency is blackholed. If the BFD tracking the IPv6 Hello adjacency times out, all IPv4 and IPv6 prefix FECs are updated.

The following behavior applies to the tracking of an mLDP FEC:

- IPv4 and IPv6 mLDP FECs are only tracked with the Hello adjacency because they do not have the concept of downstream next hop.
- The upstream LSR peer for an mLDP FEC supports the multicast upstream FRR procedures, and the upstream peer is tracked using the Hello adjacency on each link or the IPv6 transport address if there is a T-LDP session.
- The tracking of a targeted LDP peer with BFD does not change with the support of IPv6 peers. BFD tracks the transport address conveyed by the Hello adjacency that bootstrapped the LDP IPv6 session.

3.2.13 Services using SDP with an LDP IPv6 FEC

The SDP of type **ldp** with the **far-end** option using IPv6 addresses is supported. The addresses need not be of the same family (IPv6 or IPv4) for the SDP configuration to be allowed. The user can have an SDP with an IPv4 (or IPv6) control plane for the T-LDP session and an IPv6 (or IPv4) LDP FEC as the tunnel.

Because IPv6 LSP is only supported with LDP, the use of a **far-end** IPv6 address is not allowed with a BGP or RSVP/MPLS LSP. In addition, the CLI does not allow an SDP with a combination of an IPv6 LDP LSP and an IPv4 LSP of a different control plane. As a result, the following commands are blocked in the SDP configuration context when the far end is an IPv6 address:

- **bgp-tunnel**
- **lsp**
- **mixed-lsp-mode**

SDP administrative groups are not supported with an SDP using an LDP IPv6 FEC; the attempt to assign them is blocked in CLI.

Services that use LDP control plane (such as T-LDP VPLS spoke interface) have the spoke-SDP (PW) signaled with an IPv6 T-LDP session when the **far-end** option is configured to an IPv6 address. The spoke-SDP for these services binds by default to an SDP that uses a LDP IPv6 FEC, which prefix matches the far end address. In addition, the IPv6 PW control word is supported with both data plane packets and VCCV OAM packets. Hash label is also supported with the preceding services, including the signaling and negotiation of hash label support using T-LDP (Flow sub-TLV) with the LDP IPv6 control plane.

3.2.14 Mirror services

The user can configure a spoke-SDP bound to an LDP IPv6 LSP 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 one that is configured in the mirror destination service at the mirror source node. The far-end option is not supported with an IPv6 address.

3.2.14.1 Configuration at mirror source node

Use the syntax to configure at the mirror source node.

```
no spoke-sdp sdp-id:vc-id
spoke-sdp sdp-id:vc-id [create]
egress
vc-label egress-vc-label
```

The following configuration rules apply:

- The *sdp-id* must match an SDP that uses LDP IPv6 FEC.
- Configuring *egress-vc-label* is optional.

For example: **configure mirror mirror-dest 10**

3.2.14.2 Configuration at mirror destination node

Use the following syntax to configure at the mirror destination node.

```
far-end ip-address [vc-id vc-id] [ing-svc-label ingress-vc-label | tldp] [icb]
no far-end ip-address
spoke-sdp sdp-id:vc-id [create]
ingress-vc-label ingress-vc-label
exit
no shutdown
exit
exit
```

The following configuration rules apply:

- The **far-end ip-address** command is not supported with an LDP IPv6 transport tunnel. The user must reference a spoke SDP using an LDP IPv6 SDP coming from a mirror source node.
- In the **spoke-sdp sdp-id:vc-id** command, the *vc-id* value should match the **spoke-sdp** configured in the **mirror-destination** context at the mirror source node.
- Configuring *ingress-vc-label* is optional; both static and T-LDP are supported.

For example: **configure mirror mirror-dest 10 remote-source**

3.2.15 OAM Support with LDP IPv6

MPLS OAM tools **lsp-ping** and **lsp-trace** are updated to operate with LDP IPv6 and support the following:

- use of IPv6 addresses in the echo request and echo reply messages, including in DSMAP TLV, in accordance with RFC 8029
- use of LDP IPv6 prefix target FEC stack TLV, in accordance with RFC 8029
- use of IPv6 addresses in the DDMAP TLV and FEC stack change sub-TLV, in accordance with RFC 6424
- use of 127/8 IPv4 mapped IPv6 address; that is, in the range ::ffff:127/104, as the destination address of the echo request message, in accordance with RFC 8029

- use of 127/8 IPv4 mapped IPv6 address; that is, in the range ::ffff:127/104, as the **path-destination** address when the user wants to exercise a specific LDP ECMP path

The behavior at the sender and receiver nodes is updated to support both LDP IPv4 and IPv6 target FEC stack TLVs. Specifically, the following applies.

- The IP family (IPv4/IPv6) of the UDP/IP echo request message always matches the family of the LDP target FEC stack TLV the user entered in the **prefix** option.
- The **src-ip-address** option is extended to accept the IPv6 address of the sender node. If the user did not enter a source IP address, the system IPv6 address is used. If the user entered a source IP address of a different family than the LDP target FEC stack TLV, an error is returned and the test command is aborted.
- The IP family of the UDP/IP echo reply message must match that of the received echo request message.
- For **lsp-trace**, the downstream information in DSDP/DDMAP is encoded as the same family as the LDP control plane of the link LDP or targeted LDP session to the downstream peer.
- The sender node inserts a value of 69 in the Router Alert Option in the IPv6 header of the echo request packet, in accordance with RFC 5350.

Finally, **vccv-ping** and **vccv-trace** for a single-hop PW are updated to support IPv6 PW FEC 128 and FEC 129, in accordance with RFC 6829. In addition, the PW OAM control word is supported with VCCV packets when the **control-word** option is enabled on the spoke-SDP configuration. The value of the Channel Type field is set to 0x57, which indicates that the Associated Channel carries an IPv6 packet, in accordance with RFC 4385.

3.2.15.1 Configuration guidelines for LDP IPv6 OAM tools

It is recommended to ping the remote destination so IPv6 ND tables are updated on intermediate routers along the shortest path to the destination before initiating **lsp-ping** to the destination router.

3.2.16 LDP IPv6 Interoperability Considerations

3.2.16.1 Interoperability with implementations compliant with RFC 7552

The 7210 SAS implementation uses a 128-bit LSR-ID, as defined in *draft-pdutta-mpls-ldp-v2*, to establish an LDP IPv6 session with a peer LSR. This allows a routable system IPv6 address to be used by default to bring up the LDP task on the router and establish link LDP and T-LDP sessions to other LSRs, as is the common practice with LDP IPv4 in existing customer deployments. More importantly, this allows for the establishment of control plane-independent LDP IPv4 and LDP IPv6 sessions between two LSRs over the same interface or set of interfaces. The 7210 SAS implementation allows for two separate LDP IPv4 and LDP IPv6 sessions between two LSRs over the same interface or a set of interfaces because each session uses a unique LSR-ID (32-bit for IPv4 and 128-bit for IPv6).

The 7210 SAS LDP implementation does not interoperate with an implementation using a 32-bit LSR-ID, as defined in *draft-ietf-mpls-ldp-ipv6*, to establish an IPv6 LDP session. The latter specifies an LSR can send both IPv4 and IPv6 Hellos over an interface such that it can establish either an IPv4 or an IPv6 LDP session with LSRs on the same subnet. It does not allow for separate LDP IPv4 and LDP IPv6 LDP sessions between two routers.

The 7210 SAS LDP implementation should interoperate with an implementation using a 32-bit LSR-ID, as defined in *draft-ietf-mpls-ldp-ipv6*, to establish an IPv4 LDP session and to resolve both IPv4 and IPv6 prefix FECs. The 7210 SAS LDP implementation otherwise complies with all other aspects of *draft-ietf-mpls-ldp-ipv6*, including the support of the dual-stack capability TLV in the Hello message. The latter is used by an LSR to inform its peer that it is capable of establishing either an LDP IPv4 or LDP IPv6 session, and to convey the IP family preference for the LDP Hello adjacency and the resulting LDP session. This is required because the implementation described in *draft-ietf-mpls-ldp-ipv6* allows for a single session between LSRs, and both LSRs must agree if the session should be brought up using IPv4 or IPv6 when both IPv4 and IPv6 Hellos are exchanged between the two LSRs. The 7210 SAS implementation has a separate session for each IP family between two LSRs and, as such, this TLV is used to indicate the family preference and also that it supports resolving IPv6 FECs over an IPv4 LDP session.


3.2.16.2 Interoperability with implementations compliant with RFC 5036 for IPv4 LDP control plane only

The 7210 SAS implementation supports advertising and resolving IPv6 prefix FECs over an LDP IPv4 session using a 32-bit LSR-ID, in compliance with RFC 7552. When introducing an LSR based on the 7210 SAS in a LAN with a broadcast interface, it can peer with third-party LSR implementations that support RFC 7552 and LSRs that do not. When it peers, using an IPv4 LDP control plane, with a third-party LSR implementation that does not support it, the advertisement of IPv6 addresses or IPv6 FECs to that peer may cause it to bring down the IPv4 LDP session.


That is, there are deployed third-party LDP implementations that are compliant with RFC 5036 for LDP IPv4, but that are not compliant with RFC 5036 for handling IPv6 address or IPv6 FECs over an LDP IPv4 session. To resolve this issue, RFC 7552 modifies RFC 5036 by requiring implementations complying with RFC 7552 to check for the dual-stack capability TLV in the IPv4 Hello message from the peer. Without the peer advertising this TLV, an LSR must not send IPv6 addresses and FECs to that peer. The 7210 SAS implementation supports this requirement.

3.3 LDP process overview

The following figure shows the basic LDP parameter provisioning process.

Figure 39: Basic LDP parameter provisioning

The following figure shows the LDP configuration and implementation process.

Figure 40: LDP configuration and implementation

3.4 Configuring LDP with CLI

This section provides information to configure LDP using the command line interface.

3.5 LDP configuration overview

When the 7210 SAS implementation of LDP is instantiated, the protocol is in the **no shutdown** state. In addition, targeted sessions are then enabled. The default parameters for LDP are set to the documented values for targeted sessions in *draft-ietf-mpls-ldp-mib-09.txt*.

3.6 Basic LDP configuration

This chapter provides information to configure LDP and remove configuration examples of common configuration tasks.

The LDP protocol instance is created in the **no shutdown** (enabled) state.

```
A:ALU_SIM11>config>router>ldp# info
-----
    aggregate-prefix-match
        prefix-exclude "sample"
    exit
    graceful-restart
    exit
    peer-parameters
        peer 1.1.1.1
    exit
    exit
    interface-parameters
        interface "a"
    exit
    exit
    targeted-session
    exit
-----
A:ALU_SIM11>config>router>ldp#
```

3.7 Common configuration tasks

This section provides information about common configuration tasks.

3.7.1 Enabling LDP

LDP must be enabled in order for the protocol to be active. MPLS must also be enabled. MPLS is enabled in the **config>router>mpls** context.

Use the following syntax to enable LDP.

```
ldp
```

Example

```
config>router# ldp
```

Example

The following displays the enabled LDP configuration.

```
A:ALU_SIM11>config>router>ldp# info
-----
      aggregate-prefix-match
        prefix-exclude "sample"
      exit
    graceful-restart
      exit
  peer-parameters
    peer 1.1.1.1
    exit
  exit
interface-parameters
  interface "a"
  exit
exit
targeted-session
exit
-----
A:ALU_SIM11>config>router>ldp#
```

3.7.2 Configuring FEC originate parameters

A FEC can be added to the LDP IP prefix database with a specific label operation on the node. Permitted operations are pop or swap. For a swap operation, an incoming label can be swapped with a label in the range of 16 to 1048575. If a swap-label is not configured then the default value is 3.

A route table entry is required for a FEC with a pop operation to be advertised. For a FEC with a swap operation, a route-table entry must exist and user configured next-hop for swap operation must match one of the next-hops in route-table entry.

Use the following syntax to configure FEC originate parameters.

```
config>router>ldp
```

Example

```
fec-originate ip-prefix/mask [advertised-label in-label]
  next-hop ip-address [swap-label out-label]
fec-originate ip-prefix/mask [advertised-label in-label]
pop
```

Example

The following displays a FEC originate configuration example.

```
A:ALA-5>config>router# info
-----
      fec-originate 10.1.1.1/32 pop
      fec-originate 10.2.1.1/32 advertised-label 1000 next-hop 10.10.1.2
      fec-originate 10.3.1.1/32 advertised-label 1001 next-hop 10.10.2.3
swap-label 131071
session-parameters
exit
interface-parameters
```

```
        exit
        targeted-session
        exit
    exit
-----
A:ALA-5>config>router>ldp#
```

3.7.3 Configuring graceful-restart helper parameters

Graceful-restart helper advertises to its LDP neighbors by carrying the fault tolerant (FT) session TLV in the LDP initialization message, assisting the LDP in preserving its IP forwarding state across the restart. The 7210 SAS recovery is self-contained and relies on information stored internally to self-heal. This feature is only used to help third-party routers without a self-healing capability to recover.

Maximum recovery time is the time (in seconds) the sender of the TLV would like the receiver to wait, after detecting the failure of LDP communication with the sender.

Neighbor liveness time is the time (in seconds) the LSR is willing to retain its MPLS forwarding state. The time should be long enough to allow the neighboring LSRs to re-sync all the LSPs in a graceful manner, without creating congestion in the LDP control plane.

Use the following syntax to configure graceful-restart parameters.

```
config>router>ldp
[no] graceful-restart
      [no] maximum-recovery-time interval
      [no] neighbor-liveness-time interval
```

3.7.4 Applying export and import policies

Both inbound and outbound label binding filtering are supported. Inbound filtering allows a route policy to control the label bindings an LSR accepts from its peers. An import policy can accept or reject label bindings received from LDP peers.

Label bindings can be filtered based on:

- Neighbor — Match on bindings received from the specified peer.
- Interface — Match on bindings received from a neighbor or neighbors adjacent over the specified interface.
- Prefix-list — Match on bindings with the specified prefix/prefixes.

Outbound filtering allows a route policy to control the set of LDP label bindings advertised by the LSR. An export policy can control the set of LDP label bindings advertised by the router. By default, label bindings for only the system address are advertised and propagate all FECs that are received.

Matches can be based on:

- Loopback — loopback interfaces.
- All — all local subnets.
- Match — match on bindings with the specified prefix/prefixes.

Use the following syntax to apply import and export policies.

Example

```
config>router>ldp
export policy-name [policy-name...(up to 32 max)]
import policy-name [policy-name...(up to 32 max)]
```

```
A:ALU_SIM11>config>router>ldp# info
```

```
-----
          aggregate-prefix-match
            prefix-exclude "sample"
        exit
    graceful-restart
    exit
peer-parameters
    peer 1.1.1.1
    exit
exit
interface-parameters
    interface "a"
    exit
exit
targeted-session
exit
-----
```

3.7.5 Targeted session parameters

Use the following syntax to specify **targeted-session** parameters.

```
config>router# ldp
targeted-session
disable-targeted-session
hello timeout factor
keepalive timeout factor
peer ip-address
no bfd-enable
    hello timeout factor
    keepalive timeout factor
no shutdown
```

Example

The following example displays an LDP configuration example.

```
A:ALA-1>config>router>ldp# info
```

```
-----
...
          targeted-session
    hello 5000 255
    keepalive 5000 255
    peer 10.10.10.104
    no bfd-enable
    hello 2500 104
    keepalive 15 3
          exit
    exit
-----
A:ALA-1>config>router>ldp#
```


3.7.6 Interface parameters

Use the following syntax to configure interface parameters.

```
config>router# ldp
interface-parameters
hello timeout factor
keepalive timeout factor
transport-address {system|interface}
interface ip-int-name
    hello timeout factor
    keepalive timeout factor
    transport-address {system|interface}
    no shutdown
```

Example

The following example displays an interface parameter configuration example.

```
A:ALU_SIM11>config>router>ldp# info
-----
    aggregate-prefix-match
        prefix-exclude "sample"
    exit
    graceful-restart
    exit
    peer-parameters
        peer 1.1.1.1
    exit
    exit
    interface-parameters
        interface "a"
    exit
    exit
    targeted-session
    exit
-----
```

3.7.7 Peer parameters

Use the following syntax to specify interface parameters.

```
config>router# ldp
peer-parameters
peer ip-address
    auth-keychain name
    authentication-key [authentication-key|hash-key]
                        [hash|hash2]
```

```
A:ALA-1>config>router>ldp# info
-----
    peer-parameters
        peer 10.10.10.104
            authentication-key "3WErEDozxyQ" hash
    exit
    exit
```

```
        targeted-session
hello 5000 255
keepalive 5000 255
peer 10.10.10.104
no bfd-enable
hello 2500 100
keepalive 15 3
        exit
    exit
-----
A:ALA-1>config>router>ldp#
```

3.7.8 LDP signaling and services

When LDP is enabled, targeted sessions can be established to create remote adjacencies with nodes that are not directly connected. When service distribution paths (SDPs) are configured, extended discovery mechanisms enable LDP to send periodic targeted hello messages to the SDP's far-end point. The exchange of LDP hellos trigger session establishment. The SDP's signaling default enables **tldp**. The service SDP uses the targeted-session parameters configured in the **config>router>ldp>targeted-session** context.

The following example displays the command syntax usage to configure enable LDP on an MPLS SDP.

```
config>service>sdp#
signaling {off|tldp}
```

Example

The following displays an example of an SDP configuration showing the signaling default **tldp** enabled.

```
A:ALA-1>config>service>sdp# info detail
-----
description "MPLS: to-99"
far-end 10.10.10.99
lsp A_D_1
signaling tldp
path-mtu 4462
keep-alive
    hello-time 10
    hold-down-time 10
    max-drop-count 3
    timeout 5
    no message-length
    no shutdown
exit
no shutdown
-----
A:ALA-1>config>service>sdp#
```

3.8 LDP configuration management tasks

This section describes the LDP configuration management tasks.

3.8.1 Disabling LDP

The **no ldp** command disables the LDP protocol on the router. All parameters revert to the default settings. LDP must be shut down before it can be disabled.

Use the following command syntax to disable LDP:

```
no ldp
shutdown
```

3.8.2 Modifying targeted session parameters

The modification of LDP targeted session parameters does not take effect until the next time the session goes down and is re-establishes. Individual parameters cannot be deleted. The **no** form of a **targeted-session** parameter command reverts modified values back to the default.

Example

The following example displays the command syntax usage to revert targeted session parameters back to the default values.

```
config>router# ldp
config>router>ldp# targeted-session
config>router>ldp>targeted# no authentication-key
config>router>ldp>targeted# no disable-targeted-session
config>router>ldp>targeted# no hello
config>router>ldp>targeted# no keepalive
config>router>ldp>targeted# no peer 10.10.10.99
```

The following output displays the default values.

```
A:ALA-1>config>router>ldp>targeted# info detail
-----
                no disable-targeted-session
                hello 45 3
                keepalive 40 4
-----
A:ALA-1>config>router>ldp>targeted#
```

3.8.3 Modifying interface parameters

The modification of LDP targeted session parameters does not take effect until the next time the session goes down and is re-establishes. Individual parameters cannot be deleted. The **no** form of a **interface-parameter** command resets the modified values back to the defaults.

The following output displays the default values.

Example

```
A:ALU_SIM11>config>router>ldp>targ-session# info detail
-----
                no disable-targeted-session
                hello 45 3
```

```

keepalive 40 4
-----
A:ALU_SIM11>config>router>ldp>targ-session#

```

3.9 LDP command reference

3.9.1 Command hierarchies

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

3.9.1.1 LDP commands

```

config
- router
  - [no] ldp
    - [no] aggregate-prefix-match
      - prefix-exclude policy-name [policy-name...(up to 5 max)]
      - no prefix-exclude
    - [no] shutdown
    - export policy-name [policy-name...(up to 5 max)]
    - no export
    - fast-reroute
    - no fast-reroute
    - fec-originate ip-prefix/mask [advertised-label in-label] [swap-label out-label]
interface interface-name
- fec-originate ip-prefix/mask [advertised-label in-label] next-hop ip-address
[swap-label out-label]
- fec-originate ip-prefix/mask [advertised-label in-label] next-hop ip-address
[swap-label out-label] interface interface-name
- fec-originate ip-prefix/mask [advertised-label in-label] pop
- no fec-originate ip-prefix/mask interface interface-name
- no fec-originate ip-prefix/mask next-hop ip-address
- no fec-originate ip-prefix/mask next-hop ip-address interface interface-name
- no fec-originate ip-prefix/mask pop
- [no] graceful-restart
  - maximum-recovery-time interval
  - no maximum-recovery-time
  - neighbor-liveness-time interval
  - no neighbor-liveness-time
- [no] implicit-null-label
- import policy-name [policy-name...(up to 5 max)]
- interface-parameters
  - interface ip-int-name [dual-stack]
  - no interface ip-int-name
  - bfd-enable
  - no bfd-enable
  - ipv4
    - fec-type-capability
      - p2mp-ipv4 {enable | disable}
      - prefix-ipv4 {enable | disable}

```

```

- prefix-ipv6 {enable | disable}
- hello timeout factor
- no hello
- keepalive timeout factor
- no keepalive
- local-lsr-id {system | interface | interface-name interface-name}
- no local-lsr-id
- [no] shutdown
- transport-address {system | interface}
- ipv6
- fec-type-capability
- p2mp-ipv4 {enable | disable}
- prefix-ipv4 {enable | disable}
- prefix-ipv6 {enable | disable}
- hello timeout factor
- no hello
- keepalive timeout factor
- no keepalive
- local-lsr-id {system | interface}
- local-lsr-id interface-name interface-name
- no local-lsr-id
- [no] shutdown
- transport-address {system | interface}
- ipv4
- hello timeout factor
- no hello
- keepalive timeout factor
- no keepalive
- transport-address {system | interface}
- ipv6
- hello timeout factor
- no hello
- keepalive timeout factor
- no keepalive
- transport-address {system | interface}
- label-withdrawal-delay seconds
- [no] prefer-tunnel-in-tunnel
- session-parameters
- [no] peer ip-address
- [no] adv-adj-add-only
- [no] dod-label-distribution
- export-addresses policy-name [policy-name ... (up to 5 max)]
- no export-addresses
- export-prefixes policy-name [policy-name ... (up to 5 max)]
- no export-prefixes
- fec-limit limit [log-only] [threshold percentage]
- no fec-limit
- fec-type-capability
- p2mp {enable | disable}
- prefix-ipv4 {enable | disable}
- prefix-ipv6 {enable | disable}
- [no] fec129-cisco-interop
- import-prefixes policy-name [policy-name ... (up to 5 max)]
- no import-prefixes
- [no] pe-id-mac-flush-interop
- [no] shutdown
- targeted-session
- [no] disable-targeted-session
- ipv4
- hello timeout factor
- no hello
- hello-reduction {enable factor | disable}
- no hello-reduction
- keepalive timeout factor

```

```

- no keepalive
- ipv6
- hello timeout factor
- no hello
- hello-reduction {enable factor | disable}
- no hello-reduction
- keepalive timeout factor
- no keepalive
- peer ip-address
- no peer ip-address
- bfd-enable
- no bfd-enable
- hello timeout factor
- no hello
- keepalive timeout factor
- no keepalive
- local-lsr-id interface-name
- no local-lsr-id
- [no] shutdown
- no tunneling
- no lsp lsp-name
- tcp-session-parameters
- peer-transport ip-address
- auth-keychain name
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key
- [no] path-mtu-discovery
- tunnel-down-damp-time seconds
- no tunnel-down-damp-time

```

3.9.1.2 Show commands

```

show
- router
- ldp
- bindings active [fec-type prefixes] [prefix ip-prefix/mask] [egress-nh ip-address
| egress-if port-id | egress-lsp tunnel-id] [summary]
- bindings active active-ecmp
- bindings active [fec-type p2mp] [p2mp-id identifier root ip-address] [egress-
nh ip-address | egress-if port-id | egress-lsp tunnel-id] [summary]
- bindings active [fec-type p2mp] [source ip-address] group mcast-address root ip-
address]] [egress-nh ip-address] | egress-if port-id | egress-lsp tunnel-id] [summary]
- bindings active prefixes [family] [{summary | detail}] [egress-if port-id]
- bindings active prefixes [family] [{summary | detail}] [egress-lsp tunnel-id]
- bindings active prefixes [egress-nh ip-address] [family] [{summary | detail}]
- bindings active prefixes prefix ip-prefix/ip-prefix-length [{summary | detail}]
[egress-if port-id]
- bindings active prefixes prefix ip-prefix/ip-prefix-length [{summary | detail}]
[egress-lsp tunnel-id]
- bindings active prefixes prefix ip-prefix/ip-prefix-length [egress-nh ip-address]
[{summary | detail}]
- bindings fec-type {prefixes|services} [session ip-addr 4c5]] [summary| detail]
- bindings fec-type p2mp [session ip-addr[:label-space]] [summary | detail]
- bindings fec-type p2mp p2mp-id identifier root ip-address [session ip-
addr[:label-space]] [summary|detail]
- bindings fec-type p2mp root ip-address [session ip-addr[:label-space]] [summary|
detail] source ip-address group mcast-address
- bindings[fec-type fec-type [detail]] [session ip-addr[:label-space]]
- bindings [label-type] [start-label [end-label]
- bindings {prefix ip-prefix/mask [detail]] [session ip-addr[:label-space]]

```

```

- bindings prefixes prefix ip-prefix/ip-prefix-length [{summary | detail}]
[session ip-addr[:label-space]]
- bindings prefixes [family] [{summary | detail}] [session ip-addr[:label-space]]
- bindings active [prefix ip-prefix/mask]
- bindings service-id service-id [detail]
- bindings vc-type vc-type [{vc-id vc-id | agi agi} [session ip-addr[:label-
space]]]
- active
- ipv4 [summary | detail] [egress-if port-id]
- ipv4 [summary | detail] [egress-lsp tunnel-id]
- ipv4 [summary | detail] [egress-nh ip-address]
- ipv6 [summary | detail] [egress-if port-id]
- ipv6 [summary | detail] [egress-lsp tunnel-id]
- ipv6 [summary | detail] [egress-nh ip-address]
- ipv4 [session ip-addr[label-space]] [summary | detail]
- ipv6 [session ip-addr[label-space]] [summary | detail]
- discovery [{peer [ip-address]} | {interface [ip-int-name]}] [state state]
[detail]
- interface [ip-int-name | ip-address] [detail]
- parameters
- session [ip-addr[:label-space]] [detail | statistics [packet-type]]
- session-parameters [family]
- session-parameters [peer-ip-address]
- statistics
- status
- targ-peer [ip-address] [detail]
- targ-peer [detail] family
- targ-peer resource-failures [family]
- tcp-session-parameters [family]
- tcp-session-parameters [keychain keychain]
- tcp-session-parameters [transport-peer-ip-address]

```

3.9.1.3 Clear commands

```

clear
- router
- ldp
- instance
- interface [ip-int-name] [family]
- peer [ip-address] [statistics]
- session [ip-addr[:label-space]] [statistics]
- statistics

```

3.9.1.4 Debug commands

```

[no] debug
- router
- [no] ldp
- [no] interface interface-name family
- [no] event
- [no] messages
- [no] packet [detail]
- hello [detail]
- no hello
- peer ip-address
- [no] event
- [no] bindings
- [no] messages

```

```
- [no] packet
  - hello [detail]
  - no hello
  - init [detail]
  - no init
  - [no] keepalive
  - label [detail]
  - no label
```

3.9.2 Command descriptions

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

3.9.2.1 LDP configuration commands

- [Generic commands](#)
- [LDP global commands](#)
- [Session parameters commands](#)
- [Targeted session commands](#)
- [TCP session parameters commands](#)

3.9.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

config>router>ldp

config>router>ldp>targ-session>peer

config>router>ldp>interface-parameters>interface>ipv4

config>router>ldp>interface-parameters>ipv4

config>router>ldp>interface-parameters>interface>ipv6 (supported only on 7210 SAS-Mxp)

config>router>ldp>interface-parameters>ipv6 (supported only on 7210 SAS-Mxp)

config>router>ldp>aggregate-prefix-match

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 for which the default state is not indicated in the configuration file, the **shutdown** and **no shutdown** states are always indicated in system-generated configuration files.

The **no** form of this command places an entity in an administratively enabled state.

Default

no shutdown

Special Cases

LDP Protocol Handling

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

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

- The **configure router ldp** command instantiates the protocol in the **no shutdown** state and resources are allocated to enable the node to process the protocol.
- To deallocate resources, users must issue the **configure router ldp shutdown** and **configure router no ldp** commands to allow the node to boot up correctly after the reboot. It is not sufficient to issue only the **configure router ldp shutdown** command.
- See [Note](#) for more information about uninstalling LDP IPv6 IFP entries using CLI commands.

3.9.2.1.2 LDP global commands

ldp

Syntax

[no] ldp

Context

config>router

Platforms

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

Description

Commands in this context configure LDP parameters. LDP is not enabled by default and must be explicitly enabled (**no shutdown**).

To suspend the LDP protocol, use the **shutdown** command. Configuration parameters are not affected.

The **no** form of this command deletes the LDP protocol instance, removing all associated configuration parameters. The LDP instance must first be disabled using the **shutdown** command before being deleted.

aggregate-prefix-match

Syntax

[no] aggregate-prefix-match

Context

config>router>ldp

Platforms

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

Description

This command enables LDP to use the aggregate prefix match function instead of requiring an exact prefix match.

When this command is enabled, LDP performs the following procedures for all prefixes. When an LSR receives a FEC-label binding from an LDP neighbor for a specific FEC1 element, it installs the binding in the LDP FIB if:

- it is able to perform a successful longest IP match of the FEC prefix with an entry in the routing table
- the advertising LDP neighbor is the next hop to reach the FEC prefix

When the FEC-label binding has been installed in the LDP FIB, LDP programs a next-hop label forwarding entry (NHLFE) in the egress datapath to forward packets to FEC1. LDP also advertises a new FEC-label binding for FEC1 to all its LDP neighbors.

When a new prefix appears in the routing table, LDP checks the LDP FIB to determine if this prefix is a closer match for any of the installed FEC elements. If a closer match is found, LDP may have to update the NHLFE for this FEC.

When a prefix is removed from the routing table, LDP checks the LDP FIB for all FEC elements that matched this prefix to determine if another match exists in the routing table. If another match exists, it updates the NHLFE accordingly. If not, it sends a label withdraw message to its LDP neighbors to remove the binding.

If the next hop for a routing prefix changes, LDP updates the LDP FIB entry for the FEC elements that matched this prefix. It also updates the NHLFE for these FEC elements.

The **no** form of this command disables the use of the aggregate prefix match function and deletes the configuration. LDP then performs only exact prefix matching for FEC elements.

Default

no aggregate-prefix-match

prefix-exclude

Syntax

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

no prefix-exclude

Context

config>router>ldp>aggregate-prefix-match

Platforms

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

Description

This command specifies the policy name containing the prefixes to be excluded from the aggregate prefix match function. Against each excluded prefix, LDP performs an exact match of a specific FEC element prefix, instead of a longest prefix match of one or more LDP FEC element prefixes, when it receives a FEC-label binding or when a change to this prefix occurs in the routing table.

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

Default

no prefix-exclude

Parameters

policy-name

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

export

Syntax

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

no export

Context

config>router>ldp

Platforms

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

Description

This command specifies the export route policies that determine which routes are exported to LDP. Policies are configured in the **config>router>policy-options** context.

If no export policy is specified, non-LDP routes are not exported from the routing table manager to LDP, and LDP-learned routes are exported only to LDP neighbors. The current implementation of the export policy (outbound filtering) can be used only to add FECs for label propagation. The export policy does not control propagation of FECs that an LSR receives from its neighbors.

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

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.

fast-reroute

Syntax

fast-reroute

no fast-reroute

Context

config>router>ldp

Platforms

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

Description

This command enables LDP Fast-Reroute (FRR). When enabled, LDP uses both the primary next hop and LFA next hop, when available, for resolving the next hop of an LDP FEC against the corresponding prefix in the routing table. This results in LDP programming a primary NHLFE and a backup NHLFE into the forwarding engine for each next hop of a FEC prefix for the purpose of forwarding packets over the LDP FEC.

The backup NHLFE is enabled for each affected FEC next hop when any of the following events occurs.

- An LDP interface goes operationally down or is administratively shut down. In this case, LDP sends a neighbor/next-hop down message to the IOM for each LDP peer it has adjacency with over this interface.
- An LDP session to a peer goes down because the Hello or keepalive timer has expired over a specific interface. In this case, LDP sends a neighbor/next-hop down message to the IOM for this LDP peer only.
- The TCP connection used by a link LDP session to a peer goes down because, for example, next-hop tracking of the LDP transport address in RTM brings down the LDP session. In this case, LDP sends a neighbor/next-hop down message to the IOM for this LDP peer only.
- A BFD session, enabled on a T-LDP session to a peer, times out and causes the link LDP session to the same peer, which uses the same TCP connection as the T-LDP session, to also go down. In this case, LDP sends a neighbor/next-hop down message to the IOM for this LDP peer only.
- A BFD session enabled on the LDP interface to a directly connected peer times out and brings down the link LDP session to this peer. In this case, LDP sends a neighbor/next-hop down message to the IOM for this LDP peer only. BFD support on LDP interfaces is a recent feature that provides faster tracking of link LDP peers.

The **tunnel-down-dump-time** option or the **label-withdrawal-delay** option, when enabled, does not cause the corresponding timer to be activated for a FEC as long as a backup NHLFE is still available.

Because LDP can detect the loss of a neighbor/next-hop independently, it is possible that it will switch to the LFA next hop while IGP is still using the primary next hop. Also, when the interface for the previous primary next hop is restored, IGP may reconverge before LDP completes the FEC exchange with its neighbor over that interface. This may cause LDP to deprogram the LFA next hop from the FEC and blackhole traffic. To avoid this situation, IGP-LDP synchronization should be enabled on the LDP interface.

When the SPF computation determines there is more than one primary next hop for a prefix, it does not program an LFA next hop in RTM. The LDP FEC will resolve to the multiple primary next hops that provide the required protection.

The **no** form of this command disables LDP FRR.

Default

no fast-reroute

fec-originate

Syntax

fec-originate *ip-prefix/mask* [**advertised-label** *in-label*] [**swap-label** *out-label*] **interface** *interface-name*

fec-originate *ip-prefix/mask* [**advertised-label** *in-label*] **next-hop** *ip-address* [**swap-label** *out-label*]

fec-originate *ip-prefix/mask* [**advertised-label** *in-label*] **next-hop** *ip-address* [**swap-label** *out-label*]
interface *interface-name*

fec-originate *ip-prefix/mask* [**advertised-label** *in-label*] **pop**

no fec-originate *ip-prefix/mask* **interface** *interface-name*

no fec-originate *ip-prefix/mask* **next-hop** *ip-address*

no fec-originate *ip-prefix/mask* **next-hop** *ip-address* **interface** *interface-name*

no fec-originate *ip-prefix/mask* **pop**

Context

config>router>ldp

Platforms

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

Description

This command configures a way to originate a FEC (with a swap action) for which the LSR is not egress, or to originate a FEC (with a pop action) for which the LSR is egress.

Parameters

ip-prefix/mask

Specifies the information for the specified IP prefix and mask length.

Values	
<ip-address/ mask>	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:x:d.d.d.d x - [0 to FFFF]H d - [0 to 255]D ipv6-prefix-le - 0 to 128

ip-address

Specifies the IP address of the next hop of the prefix.

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:x:d.d.d.d x - [0 to FFFF]H d - [0 to 255]D

advertised-label

Keyword to specify the label advertised to the upstream peer. If not configured, the label advertised should be from the label pool. If the configured static label is not available, the IP prefix is not advertised.

out-label

Specifies the LSR to swap the label. If configured, the LSR should swap the label with the configured swap-label. If not configured, the default action is pop if the next-hop parameter is not defined.

The **next-hop**, **advertised-label**, and **swap-label** parameters are optional. If **next-hop** is configured but no **swap-label** specified, a swap occurs with label 3, such as, pop and forward to the next-hop. If the **next-hop** and **swap-label** are configured, a regular swap is performed. If no parameters are specified, a pop and route is performed.

Values 16 to 1048575

in-label

Specifies the number of labels to send to the peer associated with this FEC.

Values 32 to 1023

pop

Keyword to pop the label and transmit without the label.

interface *interface-name*

Specifies the name of the interface the label for the originated FEC is swapped to. For an unnumbered interface, this parameter is mandatory since there is no address for the next-hop. For a numbered interface, it is optional.

graceful-restart

Syntax

[no] graceful-restart

Context

config>router>ldp

Platforms

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

Description

This command enables graceful restart helper.

The **no** form of this command disables graceful restart.

Default

no graceful-restart

implicit-null-label

Syntax

[no] implicit-null-label

Context

config>router>ldp

Platforms

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

Description

This command enables the use of the implicit null label. Use this command to signal the implicit null option for all LDP FECs for which this node is the egress LER.

The **no** form of this command disables the signaling of the implicit null label.

Default

no implicit-null-label

maximum-recovery-time

Syntax

maximum-recovery-time *interval*

no maximum-recovery-time

Context

config>router>ldp

Platforms

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

Description

This command configures the local maximum recovery time.

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

Default

maximum-recovery-time 120

Parameters

interval

Specifies the length of time, in seconds.

Values 15 to 1800

neighbor-liveness-time

Syntax

neighbor-liveness-time *interval*

no neighbor-liveness-time

Context

config>router>ldp

Platforms

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

Description

This command configures the neighbor liveness time.

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

Default

neighbor-liveness-time 120

Parameters

interval

Specifies the length of time in seconds.

Values 5 to 300

import

Syntax

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

no import

Context

config>router>ldp

Platforms

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

Description

This command configures import route policies to determine which label bindings (FECs) are accepted from LDP neighbors. Policies are configured in the **config>router>policy-options** context.

If no import policy is specified, LDP accepts all label bindings from configured LDP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple import commands are issued, the last command entered 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.

label-withdrawal-delay

Syntax

label-withdrawal-delay *seconds*

Context

config>router>ldp

Platforms

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

Description

This command configures the time interval, in seconds, that LDP will delay the withdrawal of the FEC-label bindings it distributed to its neighbors when FEC is deactivated. When the timer expires, LDP sends a label withdrawal for the FEC to all its neighbors. This is applicable only to LDP transport tunnels (IPv4 prefix FECs) and is not applicable to pseudowires (service FECs).

Default

no label-withdrawal-delay

Parameters

seconds

Specifies the time that LDP delays the withdrawal of the FEC-label binding it distributed to its neighbors when FEC is deactivated.

Values 3 to 120

tunnel-down-damp-time

Syntax

tunnel-down-damp-time *seconds*

no tunnel-down-damp-time

Context

config>router>ldp

Platforms

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

Description

This command configures the time interval, in seconds, that LDP waits before posting a tunnel down event to the Tunnel Table Manager (TTM).

When LDP can no longer resolve a FEC and deactivates it, it deprograms the NHLFE in the datapath. However, it delays deleting the LDP tunnel entry in the TTM until the **tunnel-down-damp-time** timer expires. This means that users of the LDP tunnel, such as SDPs (for all services) and BGP (for Layer 3 VPNs), are not immediately notified. Traffic is still blackholed because the forwarding engine NHLFE has been deprogrammed.

If the FEC gets resolved before the **tunnel-down-damp-time** timer expires, LDP programs the forwarding engine with the new NHLFE and performs a tunnel modify event in the TTM, updating the dampened entry in the TTM with the new NHLFE information. If the FEC does not get resolved and the **tunnel-down-damp-time** timer expires, LDP posts a tunnel down event to the TTM, which deletes the LDP tunnel.

When there is an upper layer (user of LDP) that depends on the LDP control plane for failover detection, the **label-withdrawal-delay** and **tunnel-down-damp-time** options must be set to 0; for example, where a primary pseudowire does not have its own fast failover detection mechanism, and the node depends on the LDP tunnel down event to activate the standby PW.

The **no** form of this command specifies that tunnel-down events are not damped.

Parameters

seconds

Specifies the time interval, in seconds, that LDP waits before posting a tunnel down event to the TTM.

Values 0 to 20

keepalive

Syntax

keepalive *timeout factor*

no keepalive

Context

config>router>ldp>interface-parameters>interface>ipv4

config>router>ldp>interface-parameters>ipv4

config>router>ldp>interface-parameters>interface>ipv6 (supported only on 7210 SAS-Mxp)

config>router>ldp>interface-parameters>ipv6 (supported only on 7210 SAS-Mxp)

config>router>ldp>targ-session>ipv4

```
config>router>ldp>targ-session>ipv6 (supported only on 7210 SAS-Mxp)
config>router>ldp>targ-session>peer
```

Platforms

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

Description

This command configures the time interval, in seconds, that LDP waits before tearing down the session. The value of the keepalive interval is derived from the *factor* parameter.

If no LDP messages are exchanged for the configured amount of time, the LDP session is torn down. The keepalive timeout is usually three times the value of the keepalive interval. To maintain the session permanently, regardless of the activity, set the *timeout* value to zero.

When the LDP session is being set up, the keepalive timeout is negotiated to the lower of the two peers. When an operational value is agreed upon, the keepalive factor derives the value of the keepalive interval. The session must be flapped for the new settings to take effect.

The **no** form of this command at the interface level sets the *timeout* and *factor* to the values defined under the **interface-parameters** level.

The **no** form of this command at the peer level sets the *timeout* and *factor* to the values defined under the **targeted-session** level.

Default

The **keepalive timeout factor** default values, which are dependent on the CLI context, are listed in the following table.

Table 26: Keepalive timeout factor default values

Context	Timeout	Factor
config>router>ldp>if-params	30	3
config>router>ldp>targ-session	40	4
config>router>ldp>if-params>if	Inherits values from interface-parameters context	
config>router>ldp>targ-session>peer	Inherits values from targeted-session context	

Parameters

timeout

Specifies the time interval, in seconds, that LDP waits before tearing down the session.

Values 3 to 65535

factor

Specifies the number of keepalive messages, expressed as a decimal integer, that should be sent on an idle LDP session in the keepalive timeout interval.

Values 1 to 255

local-lsr-id

Syntax

local-lsr-id {system | interface | interface-name *interface-name*}

local-lsr-id interface-name *interface-name*

no local-lsr-id

Context

config>router>ldp>interface-parameters>interface>ipv4

config>router>ldp>interface-parameters>interface>ipv6 (supported only on 7210 SAS-Mxp)

Platforms

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

Description

This command configures the use of the local LDP interface address as the LSR ID to establish a link LDP adjacency and session with a directly connected LDP peer.

By default, the LDP session uses the system interface address as the LSR ID unless it is explicitly configured using this command. Although it is required to always configure the system interface on the router for the LDP protocol to come up on the node, there is no requirement to include the system interface in any routing protocol.

At initial configuration, the LDP session to the peer remains down while the interface is down.

If the user changes the LSR ID on the fly between system and interface values while the LDP session is up, LDP immediately tears down the session and attempts to re-establish it using the new LSR ID.

If the interface used as the LSR ID goes down, the LDP session goes down.

When the **interface** option is selected, the transport connection (TCP) for the link LDP session will also use the address of the local LDP interface as the transport address. If **system** is the value configured using the **config>router>ldp>interface-parameters>interface>transport-address** command, it is overridden.

The **no** form of this command returns to the default behavior of using the system interface address as the LSR ID.

Default

local-lsr-id system

Parameters

interface

Keyword to configure the local LDP interface address as the value of the LSR ID of this LDP LSR.

system

Keyword to configure the system interface address as the value of the LSR ID of this LDP LSR.

interface-name

Specifies the name of the network IP interface, up to 256 characters. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

local-lsr-id

Syntax

local-lsr-id interface-name

no local-lsr-id

Context

config>router>ldp>targeted-session>peer

Platforms

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

Description

This command enables the use of the address of a specific interface as the LSR-ID for the hello adjacency of a T-LDP session. The interface can be a regular interface or a loopback interface, including the system interface.

By default, a T-LDP session uses the system interface address as the LSR-ID. The system interface must always be configured on the router or the LDP protocol will not come up on the node. There is no requirement to include the system interface in any routing protocol though.

At initial configuration, the T-LDP session will remain down while the specified interface is down. LDP will not try to bring it up using the system interface.

If the user changes the LSR-ID on the fly while the T-LDP session is up, LDP immediately tears down the session and attempts to establish one using the new LSR-ID, regardless of operational state of the newly specified interface.

If the interface used as the LSR-ID goes down, the T-LDP session goes down.

The user-configured LSR-ID is used exclusively for extended peer discovery to establish the T-LDP hello adjacency. It is also used as the transport address for the TCP session of the LDP session when it is bootstrapped by the T-LDP hello adjacency. The user-configured LSR-ID is not used in basic peer discovery to establish a link-level LDP hello adjacency.

The **no** form of this command returns to the default behavior where the system interface address is used as the LSR-ID.

Default

no local-lsr-id

Parameters

interface-name

Specifies the name of the network IP interface, up to 32 characters. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

interface-parameters

Syntax

interface-parameters

Context

config>router>ldp

Platforms

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

Description

This command configures LDP interfaces and parameters applied to LDP interfaces.

bfd-enable

Syntax

bfd-enable

no bfd-enable

Context

config>router>ldp>targ-session

config>router>ldp>targ-session>peer

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 via the BFD command under the IP interface.

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

Default

no bfd-enable

p2mp-ipv4

Syntax

p2mp-ipv4 {enable | disable}

Context

config>router>ldp>interface-params>interface>ipv4>fec-type-capability

config>router>ldp>interface-params>interface>ipv6>fec-type-capability (supported only on 7210 SAS-Mxp)

Platforms

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

Description

This command enables and disables the IPv4 P2MP FEC capability on the interface.

Parameters

enable

Keyword to enable the IPv4 P2MP FEC capability.

disable

Keyword to disable the IPv4 P2MP FEC capability.

prefix-ipv4

Syntax

prefix-ipv4 {enable | disable}

Context

config>router>ldp>interface-params>interface>ipv4>fec-type-capability

config>router>ldp>interface-params>interface>ipv6>fec-type-capability (supported only on 7210 SAS-Mxp)

Platforms

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

Description

This command enables and disables the IPv4 prefix FEC capability on the interface.

Parameters

enable

Keyword to enable the IPv4 prefix FEC capability.

disable

Keyword to disable the IPv4 prefix FEC capability.

prefix-ipv6

Syntax

prefix-ipv6 {enable | disable}

Context

config>router>ldp>interface-params>interface>ipv4>fec-type-capability

config>router>ldp>interface-params>interface>ipv6>fec-type-capability (supported only on 7210 SAS-Mxp)

Platforms

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

Description

This command enables and disables the IPv6 prefix FEC capability on the interface.

Parameters

enable

Keyword to enable the IPv6 prefix FEC capability.

disable

Keyword to disable the IPv6 prefix FEC capability.

hello

Syntax

hello *timeout factor*

no hello

Context

config>router>ldp>interface-parameters>interface>ipv4

config>router>ldp>interface-parameters>ipv4

config>router>ldp>interface-parameters>interface>ipv6 (supported only on 7210 SAS-Mxp)

config>router>ldp>interface-parameters>ipv6 (supported only on 7210 SAS-Mxp)

config>router>ldp>targ-session>ipv4

config>router>ldp>targ-session>ipv6 (supported only on 7210 SAS-Mxp)

config>router>ldp>targ-session>peer

Platforms

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

Description

This command configures the time interval to wait before declaring a neighbor down. The *factor* parameter derives the Hello interval.

The hold time is local to the system and sent in the Hello messages to the neighbor. The hold time cannot be less than three times the hello interval.

When the LDP session is being set up, the hold down time is negotiated to the lower of the two peers. After an operational value is agreed upon, the hello factor is used to derive the value of the hello interval.

The session must be flapped for the new settings to operate.

The **no** form of this command at the **targeted-session** level sets the **hello timeout** and the **hello factor** to the default values.

The **no** form of this command at the peer level sets the **hello timeout** and the **hello factor** to the value defined under the **targeted-session** level.

Default

The following table lists the default values for the *timeout* and *factor* parameters.

Table 27: Default values for hello parameters

Context	Timeout	Factor
config>router>ldp>if-params	15	3
config>router>ldp>targ-session	45	3
config>router>ldp>if-params>if	Inherits values from the interface-parameters context	
config>router>ldp>targ-session>peer	Inherits values from the targeted-session context	

Parameters

timeout

Specifies the time interval, in seconds, that LDP waits before a neighbor goes down.

Values 3 to 65535

factor

Specifies the number of keepalive messages that should be sent on an idle LDP session in the hello timeout interval.

Values 1 to 255

hello-reduction

Syntax

```
hello-reduction {enable factor | disable}  
no hello-reduction
```

Context

```
config>router>ldp>targ-session>ipv4  
config>router>ldp>targ-session>ipv6 (supported only on 7210 SAS-Mxp)
```

Platforms

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

Description

This command enables the suppression of periodic targeted Hello messages between LDP peers after the targeted LDP session is brought up.

When this feature is enabled, the target Hello adjacency is brought up by advertising the hold-time value configured in the **hello timeout** parameter for the targeted session. The LSR node starts advertising an exponentially increasing hold-time value in the Hello message as soon as the targeted LDP session to the peer is up. Each new incremented hold-time value is sent in a number of Hello messages equal to the value of the argument factor, which represents the dampening factor, before the next exponential value is advertised. This provides time for the two peers to settle on the new value. When the hold-time reaches the maximum value of 0xffff (binary 65535), the two peers send Hello messages at a frequency of every $[(65535-1)/\text{local helloFactor}]$ seconds for the lifetime of the targeted LDP session. For example, if the local Hello factor is 3, Hello messages are sent every 21844 seconds.

The LSR node continues to compute the frequency of sending the Hello messages based on the minimum of its local hold-time value and the one advertised by its peer, as described in RFC 5036. Therefore, for the targeted LDP session to suppress the periodic Hello messages, both peers must bring their advertised hold-time to the maximum value. If one of the LDP peers does not, the frequency of the Hello messages sent by both peers continues to be governed by the smaller of the two hold-time values.

When the user enables the hello reduction option on the LSR node while the targeted LDP session to the peer is operationally up, the change takes effect immediately. That is, the LSR node starts advertising an exponentially increasing hold time value in the Hello message, starting with the current configured hold time value.

When the user disables the hello reduction option while the targeted LDP session to the peer is operationally up, the change in the hold time value from 0xffff (binary 65535) to the user-configured value for this peer takes effect immediately. The local LSR immediately advertises the value of the user-configured hold time and does not wait until the next scheduled time to send a Hello to make sure the peer adjusts its local hold timeout value immediately.

In general, any configuration change to the parameters of the T-LDP Hello adjacency (modifying the hello adjacency Hello timeout or factor, enabling or disabling hello reduction, or modifying the hello reduction factor) causes the LSR node to immediately trigger an updated Hello message with the updated hold-time value without waiting for the next scheduled time to send a Hello.

The **no** form of this command disables the hello reduction feature.

Default

no hello-reduction

Parameters

disable

Keyword to disable hello reduction.

factor

Specifies the hello reduction dampening factor.

Values 3 to 20

interface

Syntax

interface *ip-int-name* [**dual-stack**]

no interface *ip-int-name*

Context

config>router>ldp>if-params

Platforms

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

Description

This command enables LDP on the specified IP interface.

The LDP interface must be disabled using the **shutdown** command before it can be deleted.

The **no** form of this command deletes the LDP interface and all configuration information associated with the LDP interface.

Parameters

ip-int-name

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

dual-stack

Optional keyword that allows the user to explicitly indicate whether this interface should automatically create the IPv4 context. With the introduction of LDP IPv6, the creation of the interface does not automatically imply it is used for IPv4, similar to earlier IPv4 only interfaces. Therefore, the **dual-stack** keyword is an indication to the system that the user manually enables the IPv4, IPv6, or the dual-stack IPv4 and IPv6 contexts.

The following applies to the **dual-stack** keyword:

- If this keyword is configured, the IPv4 interface context is not created automatically. If it is not configured, the IPv4 interface context is created similar to the single stack LDP IPv4 interface behavior.
- This keyword is always displayed in a configuration.
- When entering an already configured interface, configuring this keyword is not required; it is ignored if configured.
- When deleting a configured interface, this keyword is not accepted in the **no** form of this command.

bfd-enable

Syntax

bfd-enable

no bfd-enable

Context

config>router>ldp>interface-parameters>interface>ipv4

Platforms

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

Description

This command enables tracking of the Hello adjacency to an LDP peer using BFD.

When this command is enabled on an LDP interface, LDP registers with BFD and starts tracking the LSR ID of all peers with which it formed Hello adjacencies over that LDP interface. The LDP hello mechanism is used to determine the remote address to be used for the BFD session. The parameters used for the BFD session, that is, **transmit-interval**, **receive-interval**, and **multiplier**, are those configured under the IP interface in existing implementation: **config>router>interface>bfd**.

When multiple links exist to the same LDP peer, a Hello adjacency is established over each link and a separate BFD session is enabled on each LDP interface. If a BFD session times out on a specific link, LDP will immediately associate the LDP session with one of the remaining Hello adjacencies and trigger the LDP FRR procedures. As soon as the last Hello adjacency goes down because of BFD timing out, the LDP session goes down and the LDP FRR procedures will be triggered.

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

Default

no bfd-enable

ipv4

Syntax

ipv4

Context

```
config>router>ldp>interface-parameters>interface
config>router>ldp>interface-parameters
config>router>ldp>targeted-session
```

Platforms

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

Description

Commands in this context configure IPv4 LDP parameters for the interface.

transport-address

Syntax

```
transport-address {interface | system}
no transport-address
```

Context

```
config>router>ldp>interface-parameters>interface>ipv4
config>router>ldp>interface-parameters>ipv4
config>router>ldp>interface-parameters>interface>ipv6 (supported only on the 7210 SAS-Mxp)
config>router>ldp>interface-parameters>ipv6 (supported only on the 7210 SAS-Mxp)
```

Platforms

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

Description

This command configures the transport address used when setting up LDP TCP sessions. The transport address, which can be set to **interface** or **system**, can be configured globally (applies to all LDP interfaces) or per interface. The most specific value is used.

With this command, you can set up the LDP interface to the connection, which can be set to the interface address or the system address. However, there can be an issue of which address to use when there are parallel adjacencies. This situation can also occur with a link and a targeted adjacency, because targeted adjacencies request the session to be set up only to the system IP address.

The **transport-address** value should not be **interface** if multiple interfaces exist between two LDP neighbors. The chosen TCP endpoint depends on the first adjacency to be formed. That is, if one LDP interface is set up as **transport-address interface** and another as **transport-address system**, the TCP endpoint addresses are determined depending on which adjacency was set up first, the TCP. After that, because the Hello message contains the LSR ID, the LDP session can be checked to verify that it is set up and match the adjacency to the session.

For any ILDP interface, as the **local-lsr-id** parameter is changed to **interface**, the **transport-address** configuration loses effectiveness because it is ignored and the ILDP sessions always uses the relevant interface IP address as transport address even though **system** is configured.

The **no** form of this command at the global level sets the transport address to the default value.

The **no** form of this command at the interface level sets the transport address to the value defined under the global level.

Default

system

Parameters

interface

Keyword to specify that the IP interface address is used to set up the LDP session between neighbors. The transport address interface cannot be used if multiple interfaces exist between two neighbors, because only one LDP session is set up between two neighbors.

system

Keyword to specify that the system IP address is used to set up the LDP session between neighbors.

prefer-tunnel-in-tunnel

Syntax

[no] prefer-tunnel-in-tunnel

Context

config>router>ldp

Platforms

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

Description

This command specifies to use tunnel-in-tunnel over a simple LDP tunnel. Specifically, the user packets for LDP FECs learned over this targeted LDP session can be sent inside an RSVP LSP, which terminates on the same egress router as the destination of the targeted LDP session. The user can specify an explicit list of RSVP LSP tunnels under the targeted LDP session or LDP will perform a lookup in the TTM for the best RSVP LSP. In the former case, only the specified LSPs will be considered to tunnel LDP user packets. In the latter case, all LSPs available to the TTM and that terminate on the same egress router as this targeted LDP session will be considered. In both cases, the metric specified under the LSP configuration is used to control this selection.

The lookup in the TTM will prefer an LDP tunnel over an LDP-over-RSVP tunnel if both are available. The tunneling operates on the data plane only. Control packets of this targeted LDP session are sent over the IGP path.

ipv6

Syntax

ipv6

Context

config>router>ldp>interface-parameters>interface

config>router>ldp>interface-parameters

config>router>ldp>targeted-session

Platforms

7210 SAS-Mxp

Description

Commands in this context configure IPv6 LDP parameters for the interface.

3.9.2.1.3 Session parameters commands

session-parameters

Syntax

session-parameters

Context

config>router>ldp

Platforms

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

Description

Commands in this context configure peer-specific parameters.

peer

Syntax

[no] peer *ip-address*

Context

config>router>ldp>session-parameters

Platforms

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

Description

This command configures parameters for an LDP peer.

Parameters

<i>ip-address</i>		
Specifies the IP address of the LDP peer 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..FFFF]H
		d - [0..255]D

adv-adj-add-only

Syntax

[no] adv-adj-addr-only

Context

config>router>ldp>session-params>peer

Platforms

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

Description

This command provides a means for an LDP router to advertise only the local IPv4 or IPv6 interfaces it uses to establish hello adjacencies with an LDP peer. By default, when a router establishes an LDP session with a peer, it advertises in an LDP Address message the addresses of all local interfaces to allow the peer to resolve LDP FECs distributed by this router. Similarly, a router sends a Withdraw Address message to all its peers to withdraw a local address if the corresponding interface went down or was deleted.

This new option reduces CPU processing when a large number of LDP neighbors come up or go down. The new CLI option is strongly recommended in mobile backhaul networks where the number of LDP peers can be large.

The **no** form of this command reverts LDP to the default behavior of advertising all local interfaces.

Default

no adv-adj-addr-only

dod-label-distribution

Syntax

[no] dod-label-distribution

Context

config>router>ldp>session-params>peer

Platforms

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

Description

This command enables the use of the LDP Downstream-on-Demand (DoD) label distribution procedures.

When this option is enabled, LDP sets the A-bit in the Label Initialization message when the LDP session to the peer is established. When both peers set the A-bit, they will both use the DoD label distribution method over the LDP session (RFC 5036).

This feature can only be enabled on a link-level LDP session and therefore applies to prefix labels only, not service labels.

As soon as the link LDP session comes up, the router sends a label request to its DoD peer for the FEC prefix corresponding to the peer LSR ID. The DoD peer LSR ID is found in the basic Hello discovery messages the peer used to establish the Hello adjacency with the router.

Similarly, if the router and the directly attached DoD peer enters into extended discovery and established a targeted LDP session, the router immediately sends a label request for the FEC prefix corresponding to the peer LSR ID found in the extended discovery messages.

However, the router will not advertise any <FEC, label> bindings, including the FEC of its own LSR-id, unless the DoD peer requested it using a Label Request Message.

When the DoD peer sends a label request for any FEC prefix, the router replies with a <FEC, label> binding for that prefix if the FEC was already activated on the router. If not, the router replies with a notification message containing the status code of "no route." The router will not attempt in the latter case to send a label request to the next-hop for the FEC prefix when the LDP session to this next-hop uses the DoD label distribution mode, therefore the reference to single-hop LDP DoD procedures.

As soon as the link LDP session comes up, the router sends a label request to its DoD peer for the FEC prefix corresponding to the peer LSR ID. The DoD peer LSR ID is found in the basic Hello discovery messages the peer used to establish the Hello adjacency with the router.

Similarly, if the router and the directly attached DoD peer enter into extended discovery and established a targeted LDP session, the router immediately sends a label request for the FEC prefix corresponding to the peer LSR ID found in the extended discovery messages. The peer address must be the peer LSR ID address.

The **no** form of this command disables the DoD label distribution with an LDP neighbor.

Default

no dod-label-distribution

export-addresses

Syntax

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

no export-addresses

Context

config>router>ldp>session-params>peer

Platforms

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

Description

This command specifies the export prefix policy to local addresses advertised to this peer.

Policies are configured in the **config>router>policy-options** context. A maximum of five policy names can be specified.

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

Default

no export-addresses

Parameters

policy-name

Specifies the export prefix route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains spaces, use double quotes to delimit the start and end of the string. The specified names must already be defined.

export-prefixes

Syntax

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

no export-prefixes

Context

config>router>ldp>session-params>peer

Platforms

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

Description

This command specifies the export route policy used to determine which prefixes received from other LDP and T-LDP peers are redistributed to this LDP peer via the LDP/T-LDP session to this peer. A prefix that is filtered out (deny) will not be exported. A prefix that is filtered in (accept) will be exported.

If no export policy is specified, all FEC prefixes learned will be exported to this LDP peer. This policy is applied in addition to the global LDP policy and targeted session policy.

Policies are configured in the **config>router>policy-options** context. A maximum of five policy names can be specified. The peer address must be the peer LSR ID address.

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

Default

no export-prefixes

Parameters

policy-name

Specifies the export prefix route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains spaces, use double quotes to delimit the start and end of the string. The specified names must already be defined.

fec-limit

Syntax

fec-limit *limit* [**log-only**] [**threshold** *percentage*]

no fec-limit

Context

config>router>ldp>session-params>peer

Platforms

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

Description

This command configures a limit on the number of FECs that an LSR will accept from a specific peer and add into the LDP label database. The limit applies to the aggregate count of all FEC types including service FEC. When the limit is reached, any FEC received will be released back to the peer. This behavior is different from the per-peer import policy, which will still accept the FEC into the label database but will not resolve it.

When the FEC limit for a peer is reached, the LSR performs the following actions:

1. generates a trap and a syslog message
2. generates an LDP notification message with the LSR overload status TLV for each LDP FEC type, including service FEC, to this peer only if this peer advertised support for the LSR overload sub-TLV via the LSR Overload Protection Capability TLV at session initialization

3. releases, with LDP Status Code of "No_Label_Resources", any new FEC, including service FEC, from this peer that exceeds the limit

If a legitimate FEC is released back to a peer while the FEC limit was exceeded, the user must have a means to replay that FEC back to the router LSR after the condition clears. This is done automatically if the peer is an SR OS-based router and supports the LDP overload status TLV (SR OS 11.0.R5 and higher). Third-party peer implementations must support the LDP overload status TLV or provide a manual command to replay the FEC.

The **threshold percentage** option allows to set a threshold value when a trap and a syslog message are generated as a warning to the user in addition to when the limit is reached. The default value for the threshold when not configured is 90%.

The **log-only** option causes a trap and syslog message to be generated when reaching the threshold and limit. However, LDP labels are not released back to the peer.

If the user decreases the limit value such that it is lower than the current number of FECs accepted from the peer, the LDP LSR raises the trap for exceeding the limit. In addition, it will set overload for peers that signaled support for the LDP overload protection capability TLV. However, no existing resolved FECs from the peer that does not support the overload protection capability TLV should be deprogrammed or released.

A different trap is released when crossing the threshold in the upward direction, when reaching the FEC limit, and when crossing the threshold in the downward direction. However the same trap will not be generated more often than two minutes apart if the number of FECs oscillates around the threshold or the FEC limit.

The **no** form of this command disables FEC limiting.

Default

no fec-limit

Parameters

limit

Specifies the aggregate count of FECs of all types that can be accepted from this LDP peer.

Values 1 to 4294967295

log-only

Keyword to enable generation of a syslog message when the threshold and limit has been reached. However, LDP labels are not released back to the peer.

percentage

Specifies the threshold value, as a percentage, that triggers a warning and syslog message and trap to be sent.

Values 1 to 100

fec-type-capability

Syntax

fec-type-capability

Context

config>router>ldp>session-params>peer

config>router>ldp>interface-params>interface>ipv4

config>router>ldp>interface-params>interface>ipv6 (supported on 7210 SAS-Mxp only)

Platforms

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

Description

Commands in this context configure FEC type capabilities for the session or interface.

p2mp

Syntax

p2mp {enable | disable}

Context

config>router>ldp>session-params>peer>fec-type-capability

Platforms

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

Description

This command enables or disables P2MP FEC capability for the session.

Default

p2mp disable

Parameters

enable

Keyword to enable P2MP FEC capability for the session.

disable

Keyword to disable P2MP FEC capability for the session.

prefix-ipv4

Syntax

prefix-ipv4 {enable | disable}

Context

config>router>ldp>session-params>peer>fec-type-capability

Platforms

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

Description

This command enables or disables the IPv4 prefix FEC capability on the session or interface.

Default

prefix-ipv4 enable

Parameters

enable

Keyword to enable IPv4 prefix FEC capability on the session or interface.

disable

Keyword to disable IPv4 prefix FEC capability on the session or interface.

prefix-ipv6

Syntax

prefix-ipv6 {enable | disable}

Context

config>router>ldp>session-params>peer>fec-type-capability

Platforms

7210 SAS-Mxp

Description

This command enables or disables the IPv6 prefix FEC capability on the session or interface.

Default

prefix-ipv4 enable

Parameters

enable

Keyword to enable IPv6 prefix FEC capability on the session or interface.

disable

Keyword to disable IPv6 prefix FEC capability on the session or interface.

fec129-cisco-interop

Syntax

[no] **fec129-cisco-interop**

Context

config>router>ldp>session-params>peer

Platforms

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

Description

This command configures whether LDP will provide translation between non-compliant FEC 129 Cisco formats. Peer LDP sessions must be manually configured toward the non-compliant Cisco PEs.

When enabled, Cisco non-compliant format is used to send and interpret received label release messages. The FEC129 SAll and TAll fields will be reversed.

The **no** form of this command disables use and support of Cisco non-compliant forms. The peer address must be the peer LSR ID address.

Default

no fec129-cisco-interop

import-prefixes

Syntax

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

no import-prefixes

Context

config>router>ldp>session-params>peer

Platforms

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

Description

This command configures the import FEC prefix policy to determine which prefixes received from this LDP peer are imported and installed by LDP on this node. If resolved, these FEC prefixes are redistributed to other LDP and T-LDP peers. A FEC prefix that is filtered out (deny) will not be imported. A FEC prefix that is filtered in (accept) will be imported.

If no import policy is specified, the node will import all prefixes received from this LDP/T-LDP peer. This policy is applied in addition to the global LDP policy and targeted session policy. Policies are configured in the **config>router>policy-options** context. A maximum of five policy names can be specified. The specified names must already be defined. The peer address must be the peer LSR ID address.

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

Default

no import-prefixes

Parameters

policy-name

Specifies the import prefix route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains spaces, use double quotes to delimit the start and end of the string.

pe-id-mac-flush-interop

Syntax

[no] pe-id-mac-flush-interop

Context

config>router>ldp>session-params>peer

Platforms

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

Description

This command enables the addition of the PE-ID TLV in the LDP MAC withdrawal (mac-flush) message, under specific conditions, and modifies the mac-flush behavior for interoperability with other vendors that do not support the flush-all-from-me vendor-specific TLV. This flag can be enabled on a per LDP peer basis and allows the flush-all-from-me interoperability with other vendors. When the pe-id-mac-flush-interop flag is enabled for a specific peer, the current mac-flush behavior is modified in terms of mac-flush generation, mac-flush propagation and behavior upon receiving a mac-flush.

The mac-flush generation will be changed depending on the type of event and according to the following rules.

- Any all-from-me mac-flush event will trigger a mac-flush all-but-mine message (RFC 4762 compliant format) with the addition of a PE-ID TLV. The PE-ID TLV contains the IP address of the sending PE.
- Any all-but-mine mac-flush event will trigger a mac-flush all-but-mine message without the addition of the PE-ID TLV, as long as the source spoke-SDP is not part of an end-point.

- Any all-but-mine mac-flush event will trigger a mac-flush all-but-mine message with the addition of the PE-ID TLV, if the source spoke-SDP is part of an end-point and the spoke-SDP goes from the down/standby state to the active state. In this case, the PE-ID TLV will contain the IP address of the PE to which the previous active spoke-SDP was connected.

Any other case will follow the existing mac-flush procedures.

When the pe-id-mac-flush-interop flag is enabled for a specific LDP peer, the mac-flush ingress processing is modified according to the following rules.

- Any received all-from-me mac-flush will follow the existing mac-flush all-from-me rules, regardless of the existence of the PE-ID.
- Any received all-but-mine mac-flush will take into account the received PE-ID, that is, all the MAC addresses associated with the PE-ID will be flushed. If the PE-ID is not included, the MAC addresses associated with the sending PE will be flushed.
- Any other case will follow the existing mac-flush procedures.

When a mac-flush message has to be propagated (for an ingress SDP-binding to an egress SDP-binding) and the pe-id-mac-flush-interop flag is enabled for the ingress and egress TLDP peers, the following behavior is observed.

- If the ingress and egress bindings are spoke-SDP, the PE will propagate the mac-flush message with its own PE-ID.
- If the ingress binding is an spoke-SDP and the egress binding a mesh-SDP, the PE will propagate the mac-flush message without modifying the PE-ID included in the PE-ID TLV.
- If the ingress binding is a mesh-SDP and the egress binding a spoke-SDP, the PE will propagate the mac-flush message with its own PE-ID.
- When ingress and egress bindings are mesh-SDP, the mac-flush message is never propagated. This is the behavior regardless of the pe-id-mac-flush-interop flag configuration.

The PE-ID TLV is never added when generating a mac-flush message on a B-VPLS if the send-bvpls-flush command is enabled in the I-VPLS. In the same way, no PE-ID is added when propagating mac-flush from a B-VPLS to a I-VPLS when the propagate-mac-flush-from-bvpls command is enabled. Mac-flush messages for peers within the same I-VPLS or within the same B-VPLS domain follow the preceding procedures.

Default

no pe-id-mac-flush-interop

3.9.2.1.4 Targeted session commands

targeted-session

Syntax

targeted-session

Context

config>router>ldp

Platforms

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

Description

This command configures targeted LDP sessions. Targeted sessions are LDP sessions between non-directly connected peers. Hello messages are sent directly to the peer platform instead of to all the routers on this subnet multicast address.

The discovery messages for an indirect LDP session are addressed to the specified peer and not to the multicast address.

disable-targeted-session

Syntax

[no] disable-targeted-session

Context

config>router>ldp>targ-session

Platforms

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

Description

This command disables support for targeted sessions. Targeted sessions are LDP sessions between non-directly connected peers. The discovery messages for an indirect LDP session are addressed to the specified peer and not to the multicast address.

The **no** form of this command enables the set up of any targeted sessions.

Default

no disable-targeted-session

peer

Syntax

[no] peer *ip-address*

Context

config>router>ldp>targeted-session

Platforms

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

Description

This command configures parameters for an LDP peer.

Parameters

ip-address

Specifies the IP address of the LDP peer in dotted-decimal notation.

tunneling

Syntax

[no] tunneling

Context

config>router>ldp>targ-session>peer

Platforms

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

Description

This command enables LDP over tunnels.

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

Default

no tunneling

lsp

Syntax

[no] lsp *lsp-name*

Context

config>router>ldp>targ-session>tunneling

Platforms

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

Description

This command configures a specific LSP destined for this peer to be used for tunneling of LDP FEC over RSVP. A maximum of four RSVP LSPs can be explicitly used for tunneling LDP FECs to the T-LDP peer.

It is not necessary to specify any RSVP LSP in this context unless there is a need to restrict the tunneling to selected LSPs. All RSVP LSPs with a to address matching that of the T-LDP peer are eligible by default.

The user can also exclude specific LSP names by using the **ldp-over-rsvp exclude** command in the **configure>router>mpls>ldp** context.

Default

no tunneling

3.9.2.1.5 TCP session parameters commands

tcp-session-parameters

Syntax

tcp-session-parameters

Context

config>router>ldp

Platforms

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

Description

Commands in this context configure parameters for the TCP transport session of an LDP session to a remote peer.

peer-transport

Syntax

peer-transport *ip-address*

no peer-transport

Context

config>router>ldp>tcp-session-parameters

Platforms

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

Description

This command configures the peer transport address, which is the destination address of the TCP connection, and not the address corresponding to the LDP LSR ID of the peer.

Default

no peer-transport

Parameters

ip-address

Specifies the IPv4 or IPv6 address of the TCP connection to the LDP peer.

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

auth-keychain

Syntax

auth-keychain *name*

Context

config>router>ldp>tcp-session-params>peer-transport

Platforms

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

Description

This command configures the TCP authentication keychain to use for the session.

Parameters

name

Specifies the name of the keychain, up to 32 characters, to use for the specified TCP session or sessions. This keychain allows the rollover of authentication keys during the lifetime of a session. The peer address must be the TCP session transport address.

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context

config>router>ldp>tcp-session-params>peer-transport

Platforms

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

Description

This command specifies the authentication key to be used between LDP peers before establishing sessions. Authentication uses the MD-5 message-based digest. The peer address must be the TCP session transport address.

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

Default

none

Parameters

authentication-key

Specifies the authentication key. The key can be any combination of ASCII characters up to 16 characters (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

Specifies the hash key. The key can be any combination of up to 33 alphanumeric characters. If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Keyword to enter the key in an encrypted form. If the **hash** keyword 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 enter the key in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assigned.

path-mtu-discovery

Syntax

[no] path-mtu-discovery

Context

config>router>ldp>tcp-session-params>peer-transport

Platforms

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

Description

This command enables path MTU discovery for the associated TCP connections. When enabled, the MTU for the associated TCP session is initially set to the egress interface MTU. The DF bit is also set so that if a router along the path of the TCP connection cannot handle a packet of a particular size without fragmenting, it sends back an ICMP message to set the path MTU for the specific session to a lower value that can be forwarded without fragmenting.

Default

no path-mtu-discovery

3.9.2.2 Show LDP commands

bindings

Syntax

```

bindings active [fec-type prefixes] [prefix ip-prefix/mask] [egress-nh ip-address | egress-if port-id |
egress-lsp tunnel-id] [summary]

bindings active active-ecmp

bindings active [fec-type p2mp] [p2mp-id identifier root ip-address] [egress-nh ip-address | egress-if
port-id | egress-lsp tunnel-id] [summary]

bindings active [fec-type p2mp] [source ip-address] group mcast-address root ip-address]] [egress-
nh ip-address] | egress-if port-id | egress-lsp tunnel-id] [summary]

bindings active prefixes [family] [{summary | detail}] [egress-if port-id]

bindings active prefixes [family] [{summary | detail}] [egress-lsp tunnel-id]

bindings active prefixes [egress-nh ip-address] [family] [{summary | detail}]

bindings prefix ip-prefix/ip-prefix-length [{summary | detail}] [egress-if port-id]

bindings prefix ip-prefix/ip-prefix-length [{summary | detail}] [egress-lsp tunnel-id]

bindings prefix ip-prefix/ip-prefix-length [egress-nh ip-address] [{summary | detail}]

bindings fec-type {prefixes|services} [session ip-addr 4c5]] [summary|detail]

bindings fec-type p2mp [session ip-addr[:label-space]] [summary|detail]

bindings fec-type p2mp p2mp-id identifier root ip-address [session ip-addr[:label-space]]
[summary|detail]

bindings p2mp-id identifier root ip-address [detail]

bindings fec-type p2mp root ip-address [session ip-addr[:label-space]] [summary|detail] source ip-
address group mcast-address

bindings [fec-type fec-type [detail]] [session ip-addr[:label-space]]

bindings label-type start-label [end-label]

bindings {prefix ip-prefix/mask [detail]} [session ip-addr[:label-space]]

bindings prefixes prefix ip-prefix/ip-prefix-length [{summary | detail}] [session ip-addr[:label-space]]

bindings prefixes [family] [{summary | detail}] [session ip-addr[:label-space]]

bindings active [prefix ip-prefix/mask]

bindings service-id service-id [detail]

bindings vc-type vc-type [{vc-id vc-id | agi agi}] [session ip-addr[:label-space]]

```

Context

```
show>router>ldp
```

Platforms

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

Description

This command displays the contents of the label information base.

Parameters

family

Specifies the family type.

Values ipv4, ipv6

fec-type fec-type

Specify the kind of FEC that the label mapping, withdraw, release and request messages are referring to.

summary

Displays information in a summarized format.

detail

Displays detailed information.

active-ecmp

Displays the LDP active bindings with ECMP routes that have been successfully installed in the hardware FIB.

session ip-addr

Displays configuration information about LDP sessions.

ip-prefix

Specify information for the specified IP prefix and mask length. Host bits must be 0.

ip-prefix-length

Specifies the length of the IP prefix.

label-space

Specifies the label space identifier that the router is advertising on the interface.

Values 0 to 65535

mask

Specifies the 32-bit address mask used to indicate the bits of an IP address that are being used for the subnet address.

Values 0 to 32

port-id

Specifies the port ID.

tunnel-id

Specifies the tunnel ID.

ip-address

Specifies the egress IP address.

start-label

Specifies a label value to begin the display.

Values 16 to 1048575

end-label

Specifies a label value to end the display.

Values 17 to 1048575

vc-type

Specifies the VC type to display.

Values ethernet, vlan, mirror

vc-id

Specifies the VC ID to display.

Values 1 to 4294967295

group multicast-address

Displays the P2MP group multicast address bindings.

Values a.b.c.d

p2mp-identifier

Displays LDP active P2MP identifier bindings.

service-id

Specifies the service ID number to display.

Values 1 to 2147483647

Output

The following outputs are examples of LDP bindings information, and [Table 28: Output fields: LDP bindings](#) describes the output fields.

Sample output

```
A:7210SAS# show router ldp bindings

=====
LDP LSR ID: 2.2.2.2
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up, D - Status Signaled Down
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        I - IES Service, R - VPRN service
        WP - Label Withdraw Pending
        BU - Alternate Next-hop for Fast Re-Route, TLV - (Type, Length: Value)
=====
LDP Prefix Bindings
=====
```

Prefix Peer	IngLbl	EgrLbl	EgrIntf/ LspId	EgrNextHop				
10.1.1.1/32 1.1.1.1	--	262143	1/1/3:12	10.11.12.1				
10.1.1.1/32 6.6.6.6	131069U	131069	--	--				
10.2.2.2/32 1.1.1.1	131071U	--	--	--				
10.2.2.2/32 6.6.6.6	131071U	--	--	--				
..... 10.6.6.6/32 6.6.6.6	--	131071	1/1/9:26	10.11.26.6				
No. of Prefix Bindings: 10								
=====								
LDP Generic P2MP Bindings								
=====								
P2MP-Id Interface	RootAddr Peer	IngLbl	EgrLbl	EgrIntf/ LspId	EgrNextHop			
8193 73732	10.1.1.1 1.1.1.1	131065U	--	--	--			
8193 73728	10.2.2.2 1.1.1.1	--	262139	1/1/3:12	10.11.12.1			
88194 73738	10.6.6.6 6.6.6.6	131054U	--	--	--			
8195 73739	10.6.6.6 6.6.6.6	131053U	--	--	--			
No. of Generic P2MP Bindings: 13								
=====								
LDP In-Band-SSM P2MP Bindings								
=====								
Source Group Interface	RootAddr Peer	IngLbl	EgrLbl	EgrIntf/	EgrNextHop			
No Matching Entries Found								
=====								
LDP Service FEC 128 Bindings								
=====								
Type	VCId	SvcId	SDPId	Peer	IngLbl	EgrLbl	LMTU	RMTU
No Matching Entries Found								
=====								
LDP Service FEC 129 Bindings								

```
=====
AGI                               SAII
Type                               TAII
      SvcId      SDPID      Peer      IngLbl  EgrLbl  LMTU  RMTU
-----
No Matching Entries Found
=====
A:7210SAS#

A:7210SAS# show router ldp bindings detail

=====
LDP P2MP Bindings
=====
-----
P2MP Type      : 1                P2MP-Id      : 100
-----
Ing Lbl        : 131065           Peer          : 10.20.1.1

Egr Lbl        : 131069   Peer          : 10.20.1.2
Egr Int/LspId  : 1/1/2   EgrNextHop  : 10.10.2.2
Egr. Flags     : None                Ing. Flags   : None

Egr Lbl        : 131067   Peer          : 10.20.1.3
Egr Int/LspId  : 1/1/3   EgrNextHop  : 10.10.3.2
Egr. Flags     : None                Ing. Flags   : None

Egr Lbl        : 131064   Peer          : 10.20.1.4
Egr Int/LspId  : 1/1/3   EgrNextHop  : 10.10.4.2
Egr. Flags     : None                Ing. Flags   : None
-----
P2MP Type      : 1                P2MP-Id      : 200
-----
Ing Lbl        : 131066           Peer          : 10.20.1.1

Egr Lbl        : 131063   Peer          : 10.20.1.2
Egr Int/LspId  : 1/1/2   EgrNextHop  : 10.10.2.2
Egr. Flags     : None                Ing. Flags   : None

Egr Lbl        : 131068   Peer          : 10.20.1.3
Egr Int/LspId  : 1/1/3   EgrNextHop  : 10.10.3.2
Egr. Flags     : None                Ing. Flags   : None
=====
A:7210SAS#

A:7210SAS# show router ldp bindings p2mp-id 8193 root 10.2.2.2 detail

=====
LDP LSR ID: 10.2.2.2
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
      WP - Label Withdraw Pending, BU - Alternate Next-hop for Fast Re-Route
=====
LDP Generic P2MP Bindings
=====
-----
P2MP Type      : 1                P2MP-Id      : 8193
```

```
-----
Root-Addr          : 10.2.2.2
-----
Ing Lbl             : --
Egr Lbl             : 262139
Egr Intf/LspId      : 1/1/3:12
EgrNextHop          : 10.11.12.1
Egr. Flags          : None
Metric              : 1
Ing. Flags           : None
Mtu                 : 1560
-----
P2MP Type           : 1
P2MP-Id             : 8193
Root-Addr           : 10.2.2.2
-----
Ing Lbl             : --
Egr Lbl             : 131059
Egr Intf/LspId      : 1/1/9:26
EgrNextHop          : 10.11.26.6
Egr. Flags          : None
Metric              : 1
Ing. Flags           : None
Mtu                 : 1560
=====
No. of Generic P2MP Bindings: 2
=====
A:7210SAS#

A:7210SAS# show router ldp bindings active fec-type p2mp

=====
LDP Generic P2MP Bindings (Active)
=====
P2MP-Id      RootAddr      IngLbl      EgrLbl      EgrIntf/      EgrNextHop
Interface      Op
-----
8193          10.1.1.1
73731         Pop          131064      --          --          --
8193          10.1.1.1
7          --
8195          10.6.6.6
73738         Pop          131058      --          --          --
-----
No. of Generic P2MP Active Bindings: 15
=====
LDP In-Band-SSM P2MP Bindings (Active)
=====
Source
Group
Interface      RootAddr      IngLbl      EgrLbl      EgrIntf/      EgrNextHop
Op
-----
No Matching Entries Found
=====
A:7210SAS#

A:7210SAS# show router ldp bindings fec-type p2mp detail
```

```

=====
LDP LSR ID: 2.2.2.2
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate Next-hop for Fast Re-Route
=====
LDP Generic P2MP Bindings
=====
-----
P2MP Type           : 1                P2MP-Id           : 8193
                          Root-Addr       : 10.1.1.1
-----
Ing Lbl             : 131053U          Peer              : 6.6.6.6
Egr Lbl             : --
Egr Int/LspId       : --
EgrNextHop          : --
Egr. Flags          : None             Ing. Flags          : None
=====
No. of Generic P2MP Bindings: 13
=====

=====
LDP In-Band-SSM P2MP Bindings
=====
No Matching Entries Found
=====
A:7210SAS#

A:7210SAS# show router ldp bindings p2mp-id 8193 root 10.2.2.2 detail

=====
LDP LSR ID: 10.2.2.2
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate Next-hop for Fast Re-Route
=====
LDP Generic P2MP Bindings
=====
-----
P2MP Type           : 1                P2MP-Id           : 8193
                          Root-Addr       : 10.2.2.2
-----
Ing Lbl             : --              Peer              : 1.1.1.1
Egr Lbl             : 262139
Egr Int/LspId       : 1/1/3:12
EgrNextHop          : 10.11.12.1
Egr. Flags          : None             Ing. Flags          : None
Metric              : 1                Mtu                : 1560
-----
P2MP Type           : 1                P2MP-Id           : 8193
                          Root-Addr       : 10.2.2.2
-----
Ing Lbl             : --              Peer              : 6.6.6.6
Egr Lbl             : 131059
Egr Int/LspId       : 1/1/9:26
EgrNextHop          : 10.11.26.6
Egr. Flags          : None             Ing. Flags          : None
Metric              : 1                Mtu                : 1560
=====
No. of Generic P2MP Bindings: 2
=====
A:7210SAS#

```

The following outputs pertain to unicast FEC resolved over an unnumbered interface.

```
A:7210SAS# # show router ldp bindings active
=====
Legend: (S) - Static          (M) - Multi-homed Secondary Support
        (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
=====
LDP Prefix Bindings (Active)
=====
Prefix          Op    IngLbl    EgrLbl    EgrIntf/LspId  EgrNextHop
-----
10.20.1.1/32    Push  --        262143    1/1/1          Unnumbered
10.20.1.1/32    Swap  262138    262143    1/1/1          Unnumbered
10.20.1.2/32    Push  --        262143    lag-1          Unnumbered
10.20.1.2/32    Swap  262139    262143    lag-1          Unnumbered
10.20.1.3/32    Pop   262143    --        --             --
10.20.1.4/32    Push  --        262143    2/1/2          Unnumbered
10.20.1.4/32    Swap  262142    262143    2/1/2          Unnumbered
10.20.1.5/32    Push  --        262143    2/1/1          Unnumbered
10.20.1.5/32    Swap  262141    262143    2/1/1          Unnumbered
10.20.1.6/32    Push  --        262140    2/1/2          Unnumbered
10.20.1.6/32    Swap  262140    262140    2/1/2          Unnumbered
-----
No. of Prefix Active Bindings: 11
=====
A:7210SAS#
A:7210SAS# show router ldp bindings
=====
LDP LSR ID: 10.20.1.3
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up, D - Status Signaled Down
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        I - IES Service, R - VPRN service
        WP - Label Withdraw Pending
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value)
=====
LDP Prefix Bindings
=====
Prefix          IngLbl    EgrLbl    EgrIntf/    EgrNextHop
Peer           LspId
-----
10.20.1.1/32    --        262143    1/1/1       Unnumbered
10.20.1.1      --
10.20.1.1/32    262138U   262142    --          --
10.20.1.2      --
10.20.1.1/32    262138U   262138    --          --
10.20.1.4      --
10.20.1.1/32    262138U   262138    --          --
10.20.1.5      --
10.20.1.2/32    262139U   262142    --          --
10.20.1.1      --
10.20.1.2/32    --        262143    lag-1       Unnumbered
10.20.1.2      --
10.20.1.2/32    262139U   262139    --          --
10.20.1.4      --
10.20.1.2/32    262139U   262139    --          --
10.20.1.5      --
10.20.1.3/32    262143U   --        --          --
10.20.1.1      --
10.20.1.3/32    262143U   --        --          --
10.20.1.2      --
10.20.1.3/32    262143U   --        --          --
10.20.1.4      --
```



```
10.20.1.3/32      262143U      --      --      --
 10.20.1.5
10.20.1.4/32      262142U      262141      --      --
 10.20.1.1
10.20.1.4/32      262142U      262141      --      --
 10.20.1.2
10.20.1.4/32      --      262143      2/1/2      Unnumbered
 10.20.1.4
10.20.1.4/32      262142U      262141      --      --
 10.20.1.5
10.20.1.5/32      262141U      262138      --      --
 10.20.1.1
10.20.1.5/32      262141U      262139      --      --
 10.20.1.2
10.20.1.5/32      262141U      262141      --      --
 10.20.1.4
10.20.1.5/32      --      262143      2/1/1      Unnumbered
 10.20.1.5
10.20.1.6/32      262140U      262140      --      --
 10.20.1.1
10.20.1.6/32      262140U      262138      --      --
 10.20.1.2
10.20.1.6/32      262140N      262140      2/1/2      Unnumbered
 10.20.1.4
10.20.1.6/32      262140U      262140      --      --
 10.20.1.5
-----
No. of Prefix Bindings: 24
=====
A:7210SAS#

A:7210SAS# show router ldp bindings detail

=====
LDP LSR ID: 10.20.1.3
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       S - Status Signaled Up, D - Status Signaled Down
       E - Epipe Service, V - VPLS Service, M - Mirror Service
       I - IES Service, R - VPRN service
       WP - Label Withdraw Pending
       BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value)
=====
LDP Prefix Bindings
=====
-----
Prefix          : 10.20.1.1/32
-----
Ing Lbl         : --                      Peer          : 10.20.1.1
Egr Lbl         : 262143
Egr Int/LspId   : 1/1/1
EgrNextHop     : Unnumbered
Egr. Flags      : None                    Ing. Flags      : None
Egr If Name     : ip-10.10.2.3
Metric         : 1000                     Mtu            : 1500
-----
Prefix          : 10.20.1.1/32
-----
Ing Lbl         : 262138U                 Peer          : 10.20.1.2
Egr Lbl         : 262142
Egr Int/LspId   : --
EgrNextHop     : --
Egr. Flags      : None                    Ing. Flags      : None
```

```

Egr If Name      : n/a
-----
Prefix           : 10.20.1.1/32
-----
Ing Lbl          : 262138U          Peer           : 10.20.1.4
Egr Lbl          : 262138
Egr Int/LspId    : --
EgrNextHop       : --
Egr. Flags       : None             Ing. Flags      : None
Egr If Name      : n/a
-----
Prefix           : 10.20.1.1/32
-----
Ing Lbl          : 262138U          Peer           : 10.20.1.5
Egr Lbl          : 262138
Egr Int/LspId    : --
EgrNextHop       : --
Egr. Flags       : None             Ing. Flags      : None
Egr If Name      : n/a
-----
Prefix           : 10.20.1.2/32
-----
Ing Lbl          : 262139U          Peer           : 10.20.1.1
Egr Lbl          : 262142
Egr Int/LspId    : --
EgrNextHop       : --
Egr. Flags       : None             Ing. Flags      : None
Egr If Name      : n/a
-----

```

A:7210SAS# show router ldp session local-addresses

=====

LDP Session Local-Addresses

=====

Session with Peer 10.20.1.2:0, Local 10.20.1.3:0

Sent Addresses:	10.1.1.1	10.10.12.3	10.10.22.3	10.20.1.3
	10.180.2.3	10.180.3.3	10.180.5.3	10.180.11.3
	10.181.2.3	10.181.3.3	10.181.5.3	10.181.11.3
	10.182.2.3	10.182.3.3	10.182.5.3	10.182.11.3

Recv Addresses:	10.2.2.2	10.10.12.2	10.20.1.2	10.180.1.2
	10.180.3.2	10.180.4.2	10.181.1.2	10.181.3.2
	10.181.4.2	10.182.1.2	10.182.3.2	10.182.4.2

Session with Peer 10.20.1.4:0, Local 10.20.1.3:0

Sent Addresses:	10.1.1.1	10.10.12.3	10.10.22.3	10.20.1.3
	10.180.2.3	10.180.3.3	10.180.5.3	10.180.11.3
	10.181.2.3	10.181.3.3	10.181.5.3	10.181.11.3
	10.182.2.3	10.182.3.3	10.182.5.3	10.182.11.3

Recv Addresses:	10.10.22.4	10.20.1.4	10.180.4.4	10.180.6.4
	10.180.9.4	10.180.11.4	10.181.4.4	10.181.6.4
	10.181.9.4	10.181.11.4	10.182.4.4	10.182.6.4
	10.182.9.4	10.182.11.4		

Session with Peer 10.20.1.5:0, Local 10.20.1.3:0

Sent Addresses:	10.1.1.1	10.10.12.3	10.10.22.3	10.20.1.3
	10.180.2.3	10.180.3.3	10.180.5.3	10.180.11.3
	10.181.2.3	10.181.3.3	10.181.5.3	10.181.11.3

```

                10.182.2.3      10.182.3.3      10.182.5.3      10.182.11.3
Recv Addresses: 10.20.1.5      10.180.5.5      10.180.6.5      10.180.10.5
                10.181.5.5      10.181.6.5      10.181.10.5     10.182.5.5
                10.182.6.5      10.182.10.5
=====
A:7210SAS#

```

The following outputs pertain to multicast P2MP FEC resolved over an unnumbered interface.

**Note:**

P2MP LSPs are only supported on 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, and 7210 SAS-T.

```

A:7210SAS# show router ldp bindings active fec-type p2mp
=====
LDP Generic P2MP Bindings (Active)
=====
P2MP-Id      RootAddr      IngLbl      EgrLbl      EgrIntf/      EgrNextHop
Interface      Op                               LspId
-----
1             10.20.1.3
73728      Push          --          262142      1/1/2:0      Unnumbered

1             10.20.1.3
73728      Push          --          262137      2/1/2:0      Unnumbered

2             10.20.1.3
73729      Push          --          262141      1/1/2:0      Unnumbered

2             10.20.1.3
73729      Push          --          262136      2/1/2:0      Unnumbered

3             10.20.1.3
73730      Push          --          262140      1/1/2:0      Unnumbered

3             10.20.1.3
73730      Push          --          262135      2/1/2:0      Unnumbered

4             10.20.1.3
73731      Push          --          262139      1/1/2:0      Unnumbered

4             10.20.1.3
73731      Push          --          262134      2/1/2:0      Unnumbered
=====
A:7210SAS#

A:7210SAS# show router ldp bindings fec-type p2m
=====
LDP LSR ID: 10.20.1.3
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Generic P2MP Bindings
=====
P2MP-Id      RootAddr      IngLbl      EgrLbl      EgrIntf/      EgrNextHop
Interface      Peer                               LspId

```

```

-----
1          10.20.1.3
73728      10.20.1.2      --      262142 1/1/2:0      Unnumbered

1          10.20.1.3
73728      10.20.1.4      --      262137 2/1/2:0      Unnumbered

2          10.20.1.3
73729      10.20.1.2      --      262141 1/1/2:0      Unnumbered

2          10.20.1.3
73729      10.20.1.4      --      262136 2/1/2:0      Unnumbered

3          10.20.1.3
73730      10.20.1.2      --      262140 1/1/2:0      Unnumbered

3          10.20.1.3
73730      10.20.1.4      --      262135 2/1/2:0      Unnumbered

4          10.20.1.3
73731      10.20.1.2      --      262139 1/1/2:0      Unnumbered

4          10.20.1.3
73731      10.20.1.4      --      262134 2/1/2:0      Unnumbered
=====
A:7210SAS#

A:7210SAS# show router ldp bindings fec-type p2mp detail
=====
LDP LSR ID: 10.20.1.3
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Generic P2MP Bindings
=====
-----
P2MP Type      : 1                      P2MP-Id      : 1
Root-Addr      : 10.20.1.3
-----
Ing Lbl        : --                      Peer        : 10.20.1.2
Egr Lbl        : 262142
Egr Int/LspId  : 1/1/2:0
EgrNextHop     : Unnumbered
Egr. Flags     : None                    Ing. Flags    : None
Egr If Name    : ip-10.180.3.3           Mtu          : 1496
Metric         : 1
-----
P2MP Type      : 1                      P2MP-Id      : 1
Root-Addr      : 10.20.1.3
-----
Ing Lbl        : --                      Peer        : 10.20.1.4
Egr Lbl        : 262137
Egr Int/LspId  : 2/1/2:0
EgrNextHop     : Unnumbered
Egr. Flags     : None                    Ing. Flags    : None
Egr If Name    : ip-10.180.11.3          Mtu          : 1496
Metric         : 1
-----
P2MP Type      : 1                      P2MP-Id      : 2
Root-Addr      : 10.20.1.3
-----
Ing Lbl        : --                      Peer        : 10.20.1.2

```

```

Egr Lbl          : 262141
Egr Int/LspId    : 1/1/2:0
EgrNextHop       : Unnumbered
Egr. Flags       : None           Ing. Flags       : None
Egr If Name      : ip-10.180.3.3
Metric           : 1              Mtu              : 1496
-----
P2MP Type        : 1              P2MP-Id       : 2
                                   Root-Addr      : 10.20.1.3
-----
Ing Lbl          : --            Peer            : 10.20.1.4
Egr Lbl          : 262136
Egr Int/LspId    : 2/1/2:0
EgrNextHop       : Unnumbered
Egr. Flags       : None           Ing. Flags       : None
Egr If Name      : ip-10.180.11.3
Metric           : 1              Mtu              : 1496
-----
A:7210SAS#

A:7210SAS# show router ldp session local-addresses
=====
LDP Session Local-Addresses
=====
-----
Session with Peer 10.20.1.2:0, Local 10.20.1.3:0
-----
Sent Addresses: 10.1.1.1      10.10.12.3    10.10.22.3    10.20.1.3
                  10.180.2.3   10.180.3.3    10.180.5.3    10.180.11.3
                  10.181.2.3   10.181.3.3    10.181.5.3    10.181.11.3
                  10.182.2.3   10.182.3.3    10.182.5.3    10.182.11.3

Recv Addresses: 10.2.2.2      10.10.12.2    10.20.1.2     10.180.1.2
                  10.180.3.2   10.180.4.2    10.181.1.2     10.181.3.2
                  10.181.4.2   10.182.1.2    10.182.3.2     10.182.4.2
-----
Session with Peer 10.20.1.4:0, Local 10.20.1.3:0
-----
Sent Addresses: 10.1.1.1      10.10.12.3    10.10.22.3    10.20.1.3
                  10.180.2.3   10.180.3.3    10.180.5.3    10.180.11.3
                  10.181.2.3   10.181.3.3    10.181.5.3    10.181.11.3
                  10.182.2.3   10.182.3.3    10.182.5.3    10.182.11.3

Recv Addresses: 10.10.22.4    10.20.1.4     10.180.4.4     10.180.6.4
                  10.180.9.4   10.180.11.4   10.181.4.4     10.181.6.4
                  10.181.9.4   10.181.11.4   10.182.4.4     10.182.6.4
                  10.182.9.4   10.182.11.4
-----
Session with Peer 10.20.1.5:0, Local 10.20.1.3:0
-----
Sent Addresses: 10.1.1.1      10.10.12.3    10.10.22.3    10.20.1.3
                  10.180.2.3   10.180.3.3    10.180.5.3    10.180.11.3
                  10.181.2.3   10.181.3.3    10.181.5.3    10.181.11.3
                  10.182.2.3   10.182.3.3    10.182.5.3    10.182.11.3

Recv Addresses: 10.20.1.5     10.180.5.5     10.180.6.5     10.180.10.5
                  10.181.5.5   10.181.6.5     10.181.10.5    10.182.5.5
                  10.182.6.5   10.182.10.5
=====
A:7210SAS#

```

Table 28: Output fields: LDP bindings

Label	Description
LDP LSR ID	Displays the LDP label switch router ID
Legend	U: Label In Use N: Label Not In Use W: Label Withdrawn S: Status Signaled Up D: Status Signaled Down E: Epipe service V: VPLS service M: Mirror service I: IES service R: VPRN service WP: Label Withdraw Pending TLV: (Type, Length: Value)
Type	Displays the service type exchanging labels; possible types displayed are VPLS, Epipe, Spoke, and Unknown
VCId	Displays the value used by each end of an SDP tunnel to identify the VC
SvcID	Displays the unique service identification number identifying the service in the service domain
Peer	Displays the IP address of the peer
P2MP-Id	Displays the P2MP ID assigned by the root to this MVPN instance
Interface	Displays the logical identifier assigned locally to identify the P2MP tunnel
RootAddr Op	Displays the IP address of the root of the P2MP tree
Op	Displays the label operation carried out (can be one of pop swap push)
EgrNextHop	Displays the next hop gateway IP address
EgrIntf/LspId	Displays the LSP tunnel ID (not the LSP path ID)
IngLbl	Displays the ingress LDP label U — Label in use

Label	Description
	R — Label released
EgrLbl	Displays the egress LDP label
LMTU	Displays the local MTU value
RMTU	Displays the remote MTU value
No. of Service Bindings	Displays the total number of LDP bindings on the router

active

Syntax

active

Context

show>router>ldp>bindings

Platforms

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

Description

Commands in this context display information about LDP active bindings.

Output

The following output is an example of LDP active bindings information, and [Table 29: Output fields: LDP active bindings](#) describes the LDP active bindings output fields.

Sample output

```
*A:Dut-A# /show router ldp bindings active

=====
LDP Bindings (IPv4 LSR ID 10.20.1.1:0)
(IPv6 LSR ID ::[0])
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
       (S) - Static      (M) - Multi-homed Secondary Support
       (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
       (C) - FEC resolved for class-based-forwarding
=====
LDP IPv4 Prefix Bindings (Active)
=====
Prefix                               Op           IngLbl      EgrLbl
EgrNextHop                           EgrIf/LspId
-----
No. of IPv4 Prefix Active Bindings: 10
```

```

=====
LDP IPv6 Prefix Bindings (Active)
=====
Prefix                               Op           IngLbl      EgrLbl
EgrNextHop                          EgrIf/LspId
-----
No Matching Entries Found
=====

LDP Generic IPv4 P2MP Bindings (Active)
=====
P2MP-Id                             Interface
RootAddr                           Op           IngLbl      EgrLbl
EgrNH                              EgrIf/LspId
-----
No Matching Entries Found
=====
*A:Dut-A#

```

Table 29: Output fields: LDP active bindings

Label	Description
LDP Bindings	Displays the LDP IPv4 and IPv6 bindings
Legend	Description of the legend used in LDP active binding output U: Label In Use N: Label Not In Use W: Label Withdrawn WP: Label Withdraw Pending BU: Alternate For Fast Re-Route (S): Static (M): Multi-homed Secondary Support (B): BGP Next Hop (BU): Alternate Next-hop for Fast Re-Route (C): FEC resolved for class-based-forwarding
LDP IPv4 Prefix Bindings (Active)	Displays the active LDP IPv4 bindings
Prefix	Displays the label binding for the LDP prefix
Op	Displays the label operation (can be either pop, swap, or push)
IngLbl	Displays the ingress LDP label
EgrLbl	Displays the egress LDP label
EgrNextHop	Displays the next-hop information for the LDP prefix

Label	Description
EgrIf/Lspld	Displays the egress interface information for the LSP tunnel ID (not the LSP path ID)
No. of IPv4 Prefix Active Bindings	Displays the number of active IPv4 prefix bindings
LDP IPv6 Prefix Bindings (Active)	Displays the active LDP IPv6 bindings
LDP Generic IPv4 P2MP Bindings (Active)	Displays the active generic LDP IPv4 P2MP bindings
P2MP-Id	Displays the P2MP ID
RootAddr	Displays the root address
EgrNH	Displays egress next-hop information

ipv4

Syntax

ipv4 [**summary** | **detail**] [**egress-if** *port-id*]
ipv4 [**summary** | **detail**] [**egress-lsp** *tunnel-id*]
ipv4 [**summary** | **detail**] [**egress-nh** *ip-address*]

Context

show>router>ldp>bindings>active

Platforms

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

Description

This command displays LDP active IPv4 bindings.

Parameters

port-id

Displays LDP active bindings by matching the egress interface.

Values *slot[/mda[/port]]*

tunnel-id

Specifies the tunnel identifier for the egress LSP.

Values 0 to 4294967295

ip-address

Displays LDP active bindings by matching the egress next hop.

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:x.d.d.d.d
		x - [0 to FFFF]H
		d - [0 to 255]D

detail

Displays detailed information about LDP active IPv4 bindings.

summary

Displays LDP active IPv4 bindings information in a summarized format.

Output

The following output is an example of LDP active IPv4 bindings information, and [Table 30: Output fields: LDP active IPv4 bindings](#) describes the LDP active IPv4 bindings output fields.

Sample output

```
*A:Dut-A# /show router ldp bindings active ipv4

=====
LDP Bindings (IPv4 LSR ID 10.20.1.1:0)
              (IPv6 LSR ID ::[0])
=====
Legend: U - Label In Use,  N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        (S) - Static          (M) - Multi-homed Secondary Support
        (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
        (C) - FEC resolved for class-based-forwarding
=====
LDP IPv4 Prefix Bindings (Active)
=====
Prefix      Op      IngLbl  EgrLbl
EgrNextHop  EgrIf/LspId
-----
10.20.1.1/32  Pop      131071  --
--           --
10.20.1.2/32  Push     --      131071
10.10.1.2     1/1/6
10.20.1.2/32  Swap     131064  131071
10.10.1.2     1/1/6
10.20.1.3/32  Push     --      131071
10.10.2.3     1/1/4
10.20.1.3/32  Swap     131063  131071
10.10.2.3     1/1/4
10.20.1.4/32  Push     --      131056
10.10.2.3     1/1/4
```

```

10.20.1.4/32          Swap      131049   131056
10.10.2.3             1/1/4

10.20.1.6/32          Push      --        131054
10.10.2.3             1/1/4

10.20.1.6/32          Swap      131048   131054
10.10.2.3             1/1/4

11.11.11.0/32(S)     Pop       131070   --
--                    --

-----
No. of IPv4 Prefix Active Bindings: 30
=====

=====
LDP Generic IPv4 P2MP Bindings (Active)
=====
P2MP-Id              Interface
RootAddr             Op          IngLbl   EgrLbl
EgrNH                EgrIf/LspId
-----
No Matching Entries Found
=====
*A:Dut-A#

```

Table 30: Output fields: LDP active IPv4 bindings

Label	Description
LDP Bindings	Displays the LDP IPv4 bindings
Legend	Description of the legend used in LDP active IPv4 binding output U: Label In Use N: Label Not In Use W: Label Withdrawn WP: Label Withdraw Pending BU: Alternate For Fast Re-Route (S): Static (M): Multi-homed Secondary Support (B): BGP Next Hop (BU): Alternate Next-hop for Fast Re-Route (C): FEC resolved for class-based-forwarding
LDP IPv4 Prefix Bindings (Active)	Displays the active LDP IPv4 bindings
Prefix	Displays the label binding for the LDP prefix
Op	Displays the label operation (can be either pop, swap, or push)

Label	Description
IngLbl	Displays the ingress LDP label
EgrLbl	Displays the egress LDP label
EgrNextHop	Displays the next-hop information for the LDP prefix
EgrIf/LspId	Displays the egress interface information for the LSP tunnel ID (not the LSP path ID)
No. of IPv4 Prefix Active Bindings	Displays the number of active IPv4 prefix bindings
LDP Generic IPv4 P2MP Bindings (Active)	Displays the active generic LDP IPv4 P2MP bindings
P2MP-Id	Displays the P2MP ID
RootAddr	Displays the root address
EgrNH	Displays egress next-hop information

ipv6

Syntax

ipv6 [**summary** | **detail**] [**egress-if** *port-id*]
ipv6 [**summary** | **detail**] [**egress-lsp** *tunnel-id*]
ipv6 [**summary** | **detail**] [**egress-nh** *ip-address*]

Context

show>router>ldp>bindings>active

Platforms

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

Description

This command displays LDP active IPv6 bindings.

Parameters

port-id

Displays LDP active bindings by matching the egress interface.

Values *slot[/mda[/port]]*

tunnel-id

Specifies the tunnel identifier for the egress LSP.

Values 0 to 4294967295

ip-address

Displays LDP active bindings by matching the egress next-hop.

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

detail

Displays detailed information about LDP active IPv6 bindings.

summary

Displays LDP active IPv6 bindings information in a summarized format.

Output

The following output is an example of LDP active IPv6 bindings information, and [Table 31: Output fields: LDP active IPv6 bindings](#) describes theLDP active IPv6 bindings output fields.

Sample output

```
*A:Dut-A# /show router ldp bindings active ipv6

=====
LDP Bindings (IPv4 LSR ID 10.20.1.1:0)
(IPv6 LSR ID ::[0])
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        (S) - Static          (M) - Multi-homed Secondary Support
        (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
        (C) - FEC resolved for class-based-forwarding
=====
LDP IPv6 Prefix Bindings (Active)
=====
Prefix                                Op          IngLbl      EgrLbl
EgrNextHop                           EgrIf/LspId
-----
No Matching Entries Found
=====
*A:Dut-A#
```

Table 31: Output fields: LDP active IPv6 bindings

Label	Description
LDP Bindings	Displays the LDP IPv6 bindings
Legend	Description of the legend used in LDP active IPv6 binding output U: Label In Use N: Label Not In Use W: Label Withdrawn WP: Label Withdraw Pending BU: Alternate For Fast Re-Route (S): Static (M): Multi-homed Secondary Support (B): BGP Next Hop (BU): Alternate Next-hop for Fast Re-Route (C): FEC resolved for class-based-forwarding
LDP IPv6 Prefix Bindings (Active)	Displays the active LDP IPv6 bindings
Prefix	Displays the label information for the IPv6 prefix
Op	Displays the label operation (can be either pop, swap, or push)
IngLbl	Displays the ingress LDP label
EgrLbl	Displays the egress LDP label
EgrNextHop	Displays the next-hop information for the LDP prefix
EgrIf/Lspld	Displays the egress interface information for the LSP tunnel ID (not the LSP path ID)

ipv4

Syntax

ipv4 [session *ip-addr*[*label-space*]] [summary | detail]

Context

show>router>ldp>bindings

Platforms

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

Description

This command displays LDP IPv4 bindings.

Parameters

ip-addr[*label-space*]
Specifies the IP address and label space identifier.

Values ipv4-address:label-space
 ipv6-address[*label-space*]
 label-space - 0 to 65535

detail
Displays detailed information about LDP IPv4 bindings.

summary
Displays LDP IPv4 bindings information in a summarized format.

Output

The following output is an example of LDP IPv4 bindings information, and [Table 32: Output fields: LDP IPv4 bindings](#) describes the LDP IPv4 bindings output fields.

Sample output

```
*A:Dut-A# /show router ldp bindings ipv4

=====
LDP Bindings (IPv4 LSR ID 10.20.1.1:0)
              (IPv6 LSR ID ::[0])
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up,  D - Status Signaled Down
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        I - IES Service, R - VPRN service
        WP - Label Withdraw Pending
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value)
=====
LDP IPv4 Prefix Bindings
=====
Prefix      IngLbl      EgrLbl
Peer      EgrIntf/LspId
EgrNextHop
-----
10.20.1.1/32      131071U      --
10.20.1.2:0      --
--

10.20.1.1/32      131071U      --
10.20.1.3:0      --
--

10.20.1.2/32      --      131071
10.20.1.2:0      1/1/6
10.10.1.2

10.20.1.2/32      131064U      131063
10.20.1.3:0      --
--
```

```

10.20.1.3/32          131063U          131057
10.20.1.2:0          --
--

10.20.1.3/32          --          131071
10.20.1.3:0          1/1/4
10.10.2.3

10.20.1.4/32          131049U          131056
10.20.1.2:0          --
--

-----
No. of IPv4 Prefix Bindings: 40
=====

=====
LDP Generic IPv4 P2MP Bindings
=====
P2MP-Id
RootAddr              Interface      IngLbl   EgrLbl
EgrNH                 EgrIf/LspId
Peer
-----
No Matching Entries Found
=====
*A:Dut-A#

```

Table 32: Output fields: LDP IPv4 bindings

Label	Description
LDP Bindings	Displays the LDP IPv4 bindings
Legend	Description of the legend used in LDP IPv4 binding output U: Label In Use N: Label Not In Use W: Label Withdrawn WP: Label Withdraw Pending BU: Alternate For Fast Re-Route (S): Static (M): Multi-homed Secondary Support (B): BGP Next Hop (BU): Alternate Next-hop for Fast Re-Route (C): FEC resolved for class-based-forwarding
LDP IPv4 Prefix Bindings	Displays the LDP IPv4 bindings
Prefix	Displays the label binding for the LDP prefix
Op	Displays the label operation (can be either pop, swap, or push)

Label	Description
IngLbl	Displays the ingress LDP label
EgrLbl	Displays the egress LDP label
EgrNextHop	Displays the next-hop information for the LDP prefix
EgrIf/Lspld	Displays the egress interface information for the LSP tunnel ID (not the LSP path ID)
No. of IPv4 Prefix Bindings	Displays the number of IPv4 prefix bindings
LDP Generic IPv4 P2MP Bindings	Displays the generic LDP IPv4 P2MP bindings
P2MP-Id	Displays the P2MP ID
RootAddr	Displays the root address
EgrNH	Displays egress next-hop information

ipv6

Syntax

ipv6 [**session** *ip-addr*[*label-space*]] [**summary** | **detail**]

Context

show>router>ldp>bindings

Platforms

7210 SAS-Mxp

Description

This command displays LDP IPv6 bindings.

Parameters

ip-addr[*label-space*]

Specifies the IP address and label space identifier.

Values ipv4-address:label-space
 ipv6-address[*label-space*]
 label-space - 0 to 65535

detail

Displays detailed information about LDP IPv6 bindings.

summary

Displays LDP IPv6 bindings information in a summarized format.

Output

The following output is an example of LDP IPv6 bindings information, and [Table 33: Output fields: LDP IPv6 bindings](#) describes the LDP IPv6 bindings output fields.

Sample output

```
*A:Dut-A# /show router ldp bindings ipv6

=====
LDP Bindings (IPv4 LSR ID 10.20.1.1:0)
              (IPv6 LSR ID ::[0])
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up, D - Status Signaled Down
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        I - IES Service, R - VPRN service
        WP - Label Withdraw Pending
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value)
=====
LDP IPv6 Prefix Bindings
=====
Prefix                               IngLbl                               EgrLbl
Peer                                 EgrIntf/LspId
EgrNextHop
-----
No Matching Entries Found
=====
*A:Dut-A#
```

Table 33: Output fields: LDP IPv6 bindings

Label	Description
LDP Bindings	Displays the LDP IPv6 bindings
Legend	Description of the legend used in LDP IPv6 binding output U: Label In Use N: Label Not In Use W: Label Withdrawn WP: Label Withdraw Pending BU: Alternate For Fast Re-Route (S): Static (M): Multi-homed Secondary Support (B): BGP Next Hop (BU): Alternate Next-hop for Fast Re-Route (C): FEC resolved for class-based-forwarding

Label	Description
LDP IPv6 Prefix Bindings	Displays the LDP IPv6 bindings
Prefix	Displays the label information for the IPv6 prefix
Op	Displays the label operation (can be either pop, swap, or push)
IngLbl	Displays the ingress LDP label
EgrLbl	Displays the egress LDP label
EgrNextHop	Displays the next-hop information for the LDP prefix
EgrIf/Lspld	Displays the egress interface information for the LSP tunnel ID (not the LSP path ID)

discovery

Syntax

discovery [{**peer** *[ip-address]*} | {**interface** *[ip-int-name]*}] [**state** *state*] [**detail**] [**adjacency-type** *type*]

Context

show>router>ldp

Platforms

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

Description

This command displays the status of the interfaces participating in LDP discovery.

Parameters

ip-address

Specifies the IP address of the peer.

ip-int-name

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

state

Specifies to display the current operational state of the adjacency.

Values established, trying, down

detail

Specifies to display detailed information.

adjacency-type type
Specifies to display the adjacency type.

Values link, targeted

Output

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

Sample output

```
*A:Dut-A# show router ldp discovery

=====
LDP IPv4 Hello Adjacencies
=====
Interface Name      Local Addr      State
AdjType            Peer Addr
-----
N/A                10.20.1.1:0    Estab
targ               10.20.1.2:0
N/A                10.20.1.1:0    Estab
targ               10.20.1.3:0
N/A                10.20.1.1:0    Estab
targ               10.20.1.7:0
a2c                10.20.1.1:0    Estab
link               10.20.1.3:0
a2b                10.20.1.1:0    Estab
link               10.20.1.2:0

-----
No. of IPv4 Hello Adjacencies: 5
=====

LDP IPv6 Hello Adjacencies
=====
Interface Name      Local Addr      State
AdjType            Peer Addr
-----
No Matching Entries Found
=====
*A:Dut-A#
```

Table 34: Output fields: LDP discovery

Label	Description
Interface Name	Displays the name of the interface
Local Addr	Displays the IP address of the originating (local) router
Peer Addr	Displays the IP address of the peer

Label	Description
Adj Type	Displays the adjacency type between the LDP peer and LDP session is targeted
State	Established — The adjacency is established Trying — The adjacency is not yet established
No. of Hello Adjacencies	Displays the total number of hello adjacencies discovered
Up Time	Displays the amount of time the adjacency has been enabled
Hold-Time Remaining	Displays the time left before a neighbor is declared to be down
Hello Mesg Recv	Displays the number of hello messages received for this adjacency
Hello Mesg Sent	Displays the number of hello messages that have been sent for this adjacency
Remote Cfg Seq No	Displays the configuration sequence number that was in the hello received when this adjacency started up. This configuration sequence number changes when there is a change of configuration.
Remote IP Address	Displays the IP address used on the remote end for the LDP session
Local Cfg Seq No	Displays the configuration sequence number that was used in the hello sent when this adjacency started up. This configuration sequence number changes when there is a change of configuration.
Local IP Address	Displays the IP address used locally for the LDP session

interface

Syntax

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

Context

show>router>ldp

Platforms

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

Description

This command displays configuration information about LDP interfaces.

Parameters

- ip-int-name**
Specifies the name of an existing interface. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.
- ip-address**
Specifies the IP address of the LDP neighbor.
- detail**
Displays detailed information.

Output

The following output is an example of LDP interface information, and [Table 35: Output fields: LDP interface](#) describes the output fields.

Sample output

```
*A:ALU_SIM11>show>router>ldp# interface
=====
LDP Interfaces
=====
Interface                Adm Opr  Hello  Hold  KA    KA    Transport
                        Factor Time  Factor Timeout Address
-----
a                        Up  Up    3      15    3      30    System
-----
No. of Interfaces: 1
=====
*A:ALU_SIM11>show>router>ldp# interface detail
*A:ALU_SIM11>show>router>ldp#
=====
LDP Interfaces (Detail)
=====
Interface "a"
-----
Admin State      : Up                Oper State      : Up
Hold Time        : 15                Hello Factor    : 3
Keepalive Timeout : 30               Keepalive Factor : 3
Transport Addr   : System            Last Modified   : 07/06/2010 10:36:59
Active Adjacencies : 1
Tunneling        : Disabled
Lsp Name         : None
=====
*A:ALU_SIM11>show>router>ldp#
```

Table 35: Output fields: LDP interface

Label	Description
Interface	Displays the interface associated with the LDP instance
Adm	Up — The LDP is administratively enabled Down — The LDP is administratively disabled
Opr	Up — The LDP is operationally enabled

Label	Description
	Down — The LDP is operationally disabled
Hello Factor	Displays the value by which the hello timeout should be divided to give the hello time, for example, the time interval, in seconds, between LDP hello messages. LDP uses hello messages to discover neighbors and to detect loss of connectivity with its neighbors.
Hold Time	Displays the hello time, also known as hold time. It is the time interval, in seconds, that LDP waits before declaring a neighbor to be down. Hello timeout is local to the system and is sent in the hello messages to a neighbor.
KA Factor	Displays the value by which the keepalive timeout should be divided to give the keepalive time, for example, the time interval, in seconds, between LDP keepalive messages. LDP keepalive messages are sent to keep the LDP session from timing out when no other LDP traffic is being sent between the neighbors.
KA Timeout	Displays the time interval, in seconds, that LDP waits before tearing down a session. If no LDP messages are exchanged during this time interval, the LDP session is torn down. Generally the value is configured to be 3 times the keepalive time (the time interval between successive LDP keepalive messages).
Auth	Enabled — Authentication using MD5 message based digest protocol is enabled Disabled — No authentication is used
No. of Interface	Displays the total number of LDP interfaces

parameters

Syntax

parameters

Context

show>router>ldp

Platforms

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

Description

This command displays configuration information about LDP parameters.

Output

The following output is an example of LDP parameter information, and [Table 36: Output fields: LDP parameters](#) describes the output fields.

Sample output

```
*A:SRU4>config>router>ldp# show router ldp parameters
=====
LDP Parameters (LSR ID 10.20.1.4)
=====
-----
Graceful Restart Parameters
-----
Nbr Liveness Time : 5 sec                      Max Recovery Time : 30
-----
Interface Parameters
-----
Keepalive Timeout : 30 sec                      Keepalive Factor : 3
Hold Time         : 15 sec                      Hello Factor      : 3
Propagate Policy  : system                     Transport Address : system
Deaggregate FECs  : False                      Route Preference  : 9
Label Distribution : downstreamUnsolicited     Label Retention   : liberal
Control Mode      : ordered                    Loop Detection    : none
-----
Targeted Session Parameters
-----
Keepalive Timeout : 40 sec                      Keepalive Factor : 4
Hold Time         : 45 sec                      Hello Factor      : 3
Passive Mode      : False                      Targeted Sessions : Enabled
=====
*A:SRU4>config>router>ldp#
```

Table 36: Output fields: LDP parameters

Label	Description
Keepalive Timeout	Displays the factor used to derive the Keepalive interval
Keepalive Factor	Displays the time interval, in seconds, that LDP waits before tearing down the session
Hold-Time	Displays the time left before a neighbor is declared to be down
Hello Factor	Displays the value by which the hello timeout should be divided to give the hello time, for example, the time interval, in seconds, between LDP hello messages. LDP uses hello messages to discover neighbors and to detect loss of connectivity with its neighbors.
Auth	Enabled — Authentication using MD5 message based digest protocol is enabled Disabled — No authentication is used
Passive-Mode	true — LDP responds only when it gets a connect request from a peer and will not attempt to actively connect to its neighbors false — LDP actively tries to connect to its peers

Label	Description
Targeted-Sessions	true — Targeted sessions are enabled false — Targeted sessions are disabled

session

Syntax

session [*ip-addr[:label-space]*] [**detail** | **statistics** [*packet-type*]] [*session-type*]

Context

show>router>ldp

Platforms

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

Description

This command displays configuration information about LDP sessions.

Parameters

ip-address

Specifies the IP address of the LDP peer.

label-space

Specifies the label space identifier that the router is advertising on the interface.

Values 0 to 65535

detail

Displays detailed information.

statistics packet-type

Specifies the packet type.

Values hello, keepalive, init, label, notification, address

session-type

Specifies to display the session type.

Values link, targeted, both

Output

The following output is an example of LDP session information, and [Table 37: Output fields: LDP session](#) describes the output fields.

Sample output

```
*A:SRU4>config>router>ldp# show router ldp session
=====
LDP Sessions
=====
Peer LDP Id      Adj Type  State      Msg Sent  Msg Recv  Up Time
-----
10.1.1.1:0       Link      Nonexistent 2         1         0d 00:00:04
10.8.100.15:0    Both      Nonexistent 14653     21054     0d 12:48:25
10.20.1.20:0     Both      Established 105187    84837     0d 12:48:27
10.20.1.22:0     Both      Established 144586    95148     0d 12:48:23
10.22.10.2:0     Link      Nonexistent 4         2         0d 00:00:16
10.22.11.2:0     Link      Nonexistent 4         4         0d 00:00:14
10.22.13.2:0     Link      Nonexistent 5         6         0d 00:00:20
10.66.33.1:0     Link      Nonexistent 6         7         0d 00:00:25
10.66.34.1:0     Link      Nonexistent 2         2         0d 00:00:05
10.66.35.1:0     Link      Nonexistent 4         4         0d 00:00:14
10.20.1.1:0      Targeted  Nonexistent 0         1         0d 00:00:04
10.20.1.3:0      Both      Established 94        97        0d 00:00:55
10.20.1.5:0      Both      Established 230866    286216    0d 12:48:27
10.20.1.110:0    Link      Nonexistent 2         2         0d 00:00:05
10.0.0.1:0       Link      Nonexistent 2         2         0d 00:00:05
-----
No. of Sessions: 15
=====

*A:SRU4>config>router>ldp#

*A:SRU4>config>router>ldp# show router ldp session 10.20.1.20:0
=====
LDP Sessions
=====
Peer LDP Id      Adj Type  State      Msg Sent  Msg Recv  Up Time
-----
10.20.1.20:0     Both      Established 105204    84859     0d 12:49:05
-----
No. of Sessions: 1
=====

*A:SRU4>config>router>ldp#

*7210SAS# show router ldp session detail
=====
LDP Sessions (Detail)
=====
Legend:  DoD - Downstream on Demand (for address FEC's only)
         DU  - Downstream Unsolicited
=====

-----
Session with Peer 10.3.3.3:0
-----
Adjacency Type   : Link                State                : Established
Up Time          : 0d 16:59:13
Max PDU Length   : 4096
Link Adjacencies : 1
Local Address    : 10.1.1.1      Peer Address         : 10.3.3.3
Local TCP Port   : 646          Peer TCP Port        : 59116
Local KA Timeout : 30           Peer KA Timeout      : 30
Mesg Sent        : 22100        Mesg Recv            : 21926
FECs Sent        : 16          FECs Recv            : 13
GR State         : Capable      Label Distribution    : DU
```

```

Nbr Liveness Time : 0          Max Recovery Time : 0
Number of Restart : 0          Last Restart Time : Never
P2MP              : Capable    MP MBB              : Capable
Dynamic Capability: Yes
Advertise         : Address

```

Session with Peer 10.2.2.2:0, Local 10.1.1.1:0

```

Adjacency Type : Both          State : Established
Up Time        : 0d 00:07:48
Max PDU Length : 4096          KA/Hold Time Remaining: 29
Link Adjacencies : 1          Targeted Adjacencies : 1
Local Address   : 10.1.1.1     Peer Address      : 10.2.2.2
Local TCP Port  : 646          Peer TCP Port     : 50980
Local KA Timeout: 30           Peer KA Timeout   : 30
Mesg Sent       : 478          Mesg Recv         : 480
FECs Sent       : 182          FECs Recv         : 170
Addrs Sent      : 13           Addrs Recv        : 16
GR State        : Capable      Label Distribution : DU
Nbr Liveness Time : 0          Max Recovery Time : 0
MP MBB          : Not Capable
Dynamic Capability: Not Capable
Advertise       : Address/Servi* BFD Operational Status: inService

```

Session with Peer 10.3.3.3:0, Local 10.1.1.1:0

```

Adjacency Type : Both          State : Established
Up Time        : 0d 00:07:48
Max PDU Length : 4096          KA/Hold Time Remaining: 29
Link Adjacencies : 1          Targeted Adjacencies : 1
Local Address   : 10.1.1.1     Peer Address      : 10.3.3.3
Local TCP Port  : 646          Peer TCP Port     : 49823
Local KA Timeout: 30           Peer KA Timeout   : 30
Mesg Sent       : 502          Mesg Recv         : 418
FECs Sent       : 124          FECs Recv         : 124
Addrs Sent      : 13           Addrs Recv        : 5
GR State        : Capable      Label Distribution : DU
Nbr Liveness Time : 0          Max Recovery Time : 0
MP MBB          : Not Capable
Dynamic Capability: Not Capable
Advertise       : Address/Servi* BFD Operational Status: inService

```

Session with Peer 10.4.4.4:0, Local 10.1.1.1:0

```

Adjacency Type : Targeted      State : Established
Up Time        : 0d 00:07:47
Max PDU Length : 4096          KA/Hold Time Remaining: 36
Link Adjacencies : 0          Targeted Adjacencies : 1
Local Address   : 10.1.1.1     Peer Address      : 10.4.4.4
Local TCP Port  : 646          Peer TCP Port     : 51307
Local KA Timeout: 40           Peer KA Timeout   : 40
Mesg Sent       : 122          Mesg Recv         : 124
FECs Sent       : 36           FECs Recv         : 36
Addrs Sent      : 13           Addrs Recv        : 3
GR State        : Capable      Label Distribution : DU
Nbr Liveness Time : 0          Max Recovery Time : 0
MP MBB          : Not Capable
Dynamic Capability: Not Capable
Advertise       : Service      BFD Operational Status: inService

```

=====

* indicates that the corresponding row element may have been truncated.

Table 37: Output fields: LDP session

Label	Description
Peer LDP ID	Displays the IP address of the LDP peer
Adj Type	Specifies that the adjacency type between the LDP peer and LDP session is targeted Link — Specifies that this adjacency is a result of a link hello Targeted — Specifies that this adjacency is a result of a targeted hello
State	Established — The adjacency is established Trying — The adjacency is not yet established
Mesg Sent	Displays the number of messages sent
Mesg Rcvd	Displays the number of messages received
Up Time	Displays the amount of time the adjacency has been enabled

session-parameters

Syntax

session-parameters [family]

session-parameters [peer-ip-address]

Context

show>router>ldp

Platforms

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

Description

This command displays LDP peer information.

Parameters

family

Specifies a peer family for which to display information.

Values ipv4, ipv6

peer-ip-address

Specifies the IP address of a targeted LDP peer for which to display information.

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

Output

The following output is an example of LDP peer information.

Sample output

```
=====
LDP IPv4 Session Parameters
=====
-----
Peer : 10.12.12.12
-----
DOD                : Disabled          Adv Adj Addr Only : Disabled
FEC129 Cisco Inter*: Disabled
PE-ID MAC Flush In*: Disabled
Fec Limit          : 0                 Fec Limit Threshold: 90
Fec Limit Log Only : Disabled
Import Policies    : None              Export Policies     : None
IPv4 Prefix Fec Cap: Enabled           IPv6 Prefix Fec Cap: Disabled
P2MP Fec Cap       : Disabled
Address Export      : None
=====
No. of IPv4 Peers: 1
=====
* indicates that the corresponding row element may have been truncated.
=====
LDP IPv6 Session Parameters
=====
No Matching Entries Found
=====
```

statistics

Syntax

statistics

Context

show>router>ldp

Platforms

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

Description

This command displays LDP statistics information.

Output

The following output is an example of LDP statistics information.

Sample output

```

=====
LDP Statistics for IPv4 LSR ID 10.20.1.1:0
                  IPv6 LSR ID 2001:db8::4444[0]
=====
Session/Discovery
-----
Active IPv4 Sess   : 5                Active IPv6 Sess   : 0
Active IPv4 LinkAdj : 21              Active IPv6 LinkAdj : 0
Active IPv4 TargAdj : 5                Active IPv6 TargAdj : 0
Active IPv4 If     : 31              Inactive IPv4 If   : 0
Active IPv6 If     : 0                Inactive IPv6 If   : 0
Active IPv4 Peers  : 5                Inactive IPv4 Peers : 0
Active IPv6 Peers  : 0                Inactive IPv6 Peers : 0
IPv4 Attempted Sess : 40              IPv6 Attempted Sess : 0
IPv4 OLoad If     : 0                IPv4 OLoad Targ Peers : 0
IPv6 OLoad If     : 0                IPv6 OLoad Targ Peers : 0
-----
Protocol Stats
-----
No Hello Err      : 0                Param Adv Err     : 0
Max PDU Err       : 0                Label Range Err   : 0
Bad LDP Id Err    : 0                Bad PDU Len Err   : 0
Bad Mesg Len Err  : 0                Bad TLV Len Err   : 0
Unknown TLV Err   : 0                Bad Proto Ver Err : 0
Malformed TLV Err : 0                Keepalive Expired Err : 9
Shutdown Notif Sent : 0              Shutdown Notif Recv : 0
-----
Prefixes
-----
IPv4 Pfx FECs Sent : 25              IPv4 Pfx FECs Recv : 25
IPv6 Pfx FECs Sent : 0                IPv6 Pfx FECs Recv : 0
IPv4PfxFec0LSessSnt : 0              IPv4PfxFec0LSessRecv : 0
IPv6PfxFec0LSessSnt : 0              IPv6PfxFec0LSessRecv : 0
IPv4PfxFecIn0Load   : 0              IPv6PfxFecIn0Load   : 0
-----
P2MP
-----
IPv4 P2MP FECs Sent : 0                IPv4 P2MP FECs Recv : 0
IPv6 P2MP FECs Sent : 0                IPv6 P2MP FECs Recv : 0
IPv4P2MPFec0LSessSn : 0              IPv4P2MPFec0LSessRecv : 0
IPv6P2MPFec0LSessSn : 0              IPv6P2MPFec0LSessRecv : 0
IPv4P2MPFecIn0Load  : 0              IPv6P2MPFecIn0Load  : 0
-----
Services
-----
Svc FEC128s Sent   : 0                Svc FEC128s Recv   : 0
Svc FEC129s Sent   : 0                Svc FEC129s Recv   : 0
Svc Fec128 0LSessSn : 0              Svc Fec128 0LSessRecv : 0
Svc Fec129 0LSessSn : 0              Svc Fec129 0LSessRecv : 0
Svc Fec128 In0Load  : 0              Svc Fec129 In0Load  : 0
=====

```

status

Syntax
status

Context
show>router>ldp

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

Description
This command displays LDP status information.

Output
The following output is an example of LDP status information, and [Table 38: Output fields: LDP status](#) describes the output fields.

Sample output

```
A:7210SAS# show router ldp status

=====
LDP Status for LSR ID 10.1.1.1
=====
Admin State      : Up                Oper State      : Up
Created at       : 11/30/2009 21:22:31 Up Time         : 0d 16:45:41
Oper Down Reason : n/a              Oper Down Events : 1
Last Change      : 11/30/2009 21:22:31 Tunn Down Damp Time : 20 sec
Label Withdraw Del*: 120 sec         Implicit Null Label : Disabled
Short. TTL Prop Lo*: Enabled         Short. TTL Prop Tran*: Enabled
Import Policies  :                   Export Policies   :
None                                                    ldp-reject
Aggregate Prefix : True              Agg Prefix Policies: None
Dynamic Capability : True            P2MP Capable      : True
MP MBB Capable    : True            MP MBB Timer       : 3
Active Adjacencies : 3               Active Sessions    : 2
Active Interfaces  : 23              Inactive Interfaces : 8
Active Peers       : 4               Inactive Peers      : 1
Addr FECs Sent     : 34              Addr FECs Recv      : 15
Serv FECs Sent     : 0               Serv FECs Recv      : 0
P2MP FECs Sent     : 0               P2MP FECs Recv      : 0
Attempted Sessions : 0
No Hello Err       : 0               Param Adv Err       : 0
Max PDU Err        : 0               Label Range Err     : 0
Bad LDP Id Err     : 4296            Bad PDU Len Err     : 0
Bad Mesg Len Err   : 0               Bad TLV Len Err     : 0
Unknown TLV Err    : 0
Malformed TLV Err  : 0               Keepalive Expired Err: 0
Shutdown Notif Sent: 0               Shutdown Notif Recv  : 0
No Matching Entries Found

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

Table 38: Output fields: LDP status

Label	Description
Admin State	Up — The LDP is administratively enabled Down — The LDP is administratively disabled
Oper State	Up — The LDP is operationally enabled Down — The LDP is operationally disabled
Created at	Displays the date and time when the LDP instance was created
Up Time	Displays the time, in hundredths of seconds, that the LDP instance has been operationally up
Last Change	Displays the date and time when the LDP instance was last modified
Oper Down Events	Displays the number of times the LDP instance has gone operationally down since the instance was created
Active Adjacencies	Displays the number of active adjacencies (established sessions) associated with the LDP instance
Active Sessions	Displays the number of active sessions (session in some form of creation) associated with the LDP instance
Active Interfaces	Displays the number of active (operationally up) interfaces associated with the LDP instance
Inactive Interfaces	Displays the number of inactive (operationally down) interfaces associated with the LDP instance
Active Peers	Displays the number of active LDP peers
Inactive Peers	Displays the number of inactive LDP peers
Addr FECs Sent	Displays the number of labels that have been sent to the peer associated with this FEC
Addr FECs Recv	Displays the number of labels that have been received from the peer associated with this FEC
Serv FECs Sent	Displays the number of labels that have been sent to the peer associated with this FEC
Serv FECs Recv	Displays the number of labels that have been received from the peer associated with this FEC
Attempted Sessions	Displays the total number of attempted sessions for this LDP instance

Label	Description
No Hello Err	Displays the total number of "Session Rejected" or "No Hello Error" notification messages sent or received by this LDP instance
Param Adv Err	Displays the total number of "Session Rejected" or "Parameters Advertisement Mode Error" notification messages sent or received by this LDP instance
Max PDU Err	Displays the total number of "Session Rejected" or "Parameters Max PDU Length Error" notification messages sent or received by this LDP instance
Label Range Err	Displays the total number of "Session Rejected" or "Parameters Label Range Error" notification messages sent or received by this LDP instance
Bad LDP Id Err	Displays the number of bad LDP identifier fatal errors detected for sessions associated with this LDP instance
Bad PDU Len Err	Displays the number of bad PDU length fatal errors detected for sessions associated with this LDP instance
Bad Mesg Len Err	Displays the number of bad message length fatal errors detected for sessions associated with this LDP instance
Bad TLV Len Err	Displays the number of bad TLV length fatal errors detected for sessions associated with this LDP instance
Malformed TLV Err	Displays the number of malformed TLV value fatal errors detected for sessions associated with this LDP instance
Shutdown Notif Sent	Displays the number of shutdown notifications sent related to sessions associated with this LDP instance
Keepalive Expired Err	Displays the number of session Keepalive timer expired errors detected for sessions associated with this LDP instance
Shutdown Notif Recv	Displays the number of shutdown notifications received related to sessions associated with this LDP instance

targ-peer

Syntax

targ-peer [*ip-address*] [*detail*]

targ-peer [*detail*] *family*

targ-peer resource-failures [*family*]

Context

show>router>ldp

Platforms

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

Description

This command displays targeted LDP peer information.

Parameters

detail

Displays detailed information.

family

Specifies a peer family for which to display information.

Values ipv4, ipv6

ip-address

Specifies the IP address of a targeted LDP peer for which to display information.

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

resource-failures

Displays resource failure information for targeted LDP peers.

Output

The following output is an example of targeted LDP peer information, and [Table 39: Output fields: LDP targeted peers](#) describes the output fields.

Sample output

=====						
LDP IPv4 Targeted Peers						
=====						
Peer	Adm/ Opr	Hello Fctr	Hold Time	KA Fctr	KA Time	Auto Created

10.1.1.1	Up/Up	3	45	4	40	yes
10.2.2.2	Up/Up	3	45	4	40	yes
10.3.3.3	Up/Up	3	45	4	40	yes

```

10.5.5.5                Up/Up  3    45    4    40    yes
10.6.6.6                Up/Up  3    45    4    40    yes

-----
No. of IPv4 Targeted Peers: 5
=====

=====
LDP IPv6 Targeted Peers
=====
Peer                    Adm/   Hello Hold  KA   KA   Auto
                      Opr    Fctr Time Fctr Time Created
-----
No Matching Entries Found
=====

```

Table 39: Output fields: LDP targeted peers

Label	Description
Peer	Displays the IP address of the peer
Adm	Up — indicates that LDP is administratively enabled Down — indicates that LDP is administratively disabled
Opr	Up — indicates that LDP is operationally enabled Down — indicates that LDP is operationally disabled
Hello Factor	The value by which the hello timeout should be divided to give the hello time; that is, the time interval, in seconds, between LDP Hello messages. LDP uses hello messages to discover neighbors and to detect loss of connectivity with its neighbors.
Hold Time	Displays the time interval, in seconds, that LDP waits before declaring a neighbor to be down. Hold time (also known as Hello time) is local to the system and is sent in the hello messages to a neighbor.
Keepalive Factor	Displays the value by which the keepalive timeout should be divided to give the keepalive time; that is, the time interval, in seconds, between LDP keepalive messages. LDP keepalive messages are sent to keep the LDP session from timing out when no other LDP traffic is being sent between the neighbors.
Keepalive Timeout	Displays the time interval, in seconds, that LDP waits before tearing down a session. If no LDP messages are exchanged during this time interval, the LDP session is torn down. Generally the value is configured to be three times the keepalive time (the time interval between successive LDP keepalive messages).

Label	Description
Auto Create	Specifies whether a targeted peer was automatically created through a Service Manager. For an LDP interface, this value is always false.
No. of Peers	Displays the total number of LDP peers

tcp-session-parameters

Syntax

```
tcp-session-parameters [family]
tcp-session-parameters [keychain keychain]
tcp-session-parameters [transport-peer-ip-address]
```

Context

```
show>router>ldp
```

Platforms

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

Description

This command displays information about the TCP transport session of an LDP peer.

Parameters

family

Specifies a peer family for which to display information.

Values ipv4, ipv6

keychain

Specifies the name of an auth-keychain, up to 32 characters.

transport-peer-ip-address

Specifies the IP address of a TCP transport peer for which to display information.

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

3.9.2.3 Clear commands

instance

Syntax

instance

Context

clear>router>ldp

Platforms

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

Description

This command resets the LDP instance.

interface

Syntax

interface [*ip-int-name*] [**family**]

Context

clear>router>ldp

Platforms

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

Description

This command restarts or clears statistics for LDP interfaces.

Parameters

ip-int-name

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

family

Specifies the family type.

Values ipv4, ipv6

peer

Syntax

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

Context

clear>router>ldp

Platforms

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

Description

This command restarts or clears statistics for LDP targeted peers.

Parameters

ip-address

Specifies the IP address of a targeted peer.

statistics

Clears only the statistics for a targeted peer

session

Syntax

session [*ip-addr[:label-space]*] [**statistics**]

Context

clear>router>ldp

Platforms

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

Description

This command restarts or clears statistics for LDP sessions.

Parameters

label-space

Specifies the label space identifier that the router is advertising on the interface.

Values 0 to 65535

statistics

Clears only the statistics for a session.

statistics

Syntax

statistics

Context

clear>router>ldp

Platforms

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

Description

This command clears LDP instance statistics.

3.9.2.4 Debug commands

ldp

Syntax

[no] ldp

Context

debug>router

Platforms

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

Description

This command configures LDP debugging.

interface

Syntax

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

Context

debug>router>ldp

Platforms

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

Description

This command debugs an LDP interface.

Parameters

interface-name

Specifies the name of an existing interface.

family

Specifies the family type.

Values ipv4, ipv6

peer

Syntax

[no] peer *ip-address*

Context

debug>router>ldp

Platforms

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

Description

This command debugs an LDP peer.

Parameters

ip-address

Specifies the IP address of the LDP peer.

event

Syntax

[no] event

Context

debug>router>ldp>peer

Platforms

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

Description

This command configures debugging for specific LDP events.

bindings

Syntax

[no] bindings

Context

debug>router>ldp>peer>event

Platforms

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

Description

This command displays debugging information about addresses and label bindings learned from LDP peers for LDP bindings.

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

messages

Syntax

[no] messages

Context

debug>router>ldp>peer>event

Platforms

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

Description

This command displays specific information (for example, message type, source, and destination) about LDP messages sent to and received from LDP peers.

The **no** form of this command disables debugging output for LDP messages.

packet

Syntax

packet [detail]

no packet

Context

debug>router>ldp>peer

Platforms

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

Description

This command enables debugging for specific LDP packets.

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

Parameters

detail

Keyword to display detailed information.

hello

Syntax

hello [**detail**]

no hello

Context

debug>router>ldp>peer>packet

Platforms

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

Description

This command enables debugging for LDP hello packets.

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

Parameters

detail

Keyword to display detailed information.

init

Syntax

init [**detail**]

no init

Context

debug>router>ldp>peer>packet

Platforms

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

Description

This command enables debugging for LDP Init packets.

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

Parameters

detail

Displays detailed information.

keepalive

Syntax

[no] keepalive

Context

debug>router>ldp>peer>packet

Platforms

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

Description

This command enables debugging for LDP Keepalive packets.

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

label

Syntax

label [detail]

no label

Context

debug>router>ldp>peer>packet

Platforms

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

Description

This command enables debugging for LDP Label packets.

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

Parameters

detail

Keyword to display detailed information.

4 PCEP

This chapter provides information about the Path Computation Element (PCE) Communication Protocol (PCEP).



Note:

PCEP is only supported on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

4.1 Introduction to PCEP



Note:

The 7210 SAS operates as a PCE Client (PCC) only, supporting PCC capabilities for RSVP-TE LSPs. References to PCE router operation apply to the Network Services Platform (NSP) or to a virtualized Service Router (VSR) operating in the control and management domain of the NSP, and are included for informational purposes only.

PCEP is one of several protocols used for communication between a wide area network (WAN) software-defined network (SDN) controller and network elements.

The Nokia WAN SDN Controller is known as the NSP. The NSP is a set of applications built on a common framework that hosts and integrates them by providing common functions. The applications are developed in a Java environment.


The NSP provides two major functions:

- programmable multi-vendor service provisioning
- network resource control, including resource management at Layer 0 (optical path), Layer 1 (ODU path), Layer 2 (MPLS tunnel), and at the IP flow level

The network discovery and control function implements a common set of standards-based southbound interfaces to the network elements for both topology discovery and tunnel and flow programming. A virtual SR OS (vSROS) applies the southbound interfaces to the network elements and the adaptation layer to the applications. The southbound interfaces include IGP and the Network Functions Manager - Packet (NFM-P) for topology discovery, PCEP for handling path computation requests and LSP state updates with the network elements, and forwarding plane programming protocols such as Openflow, BGP flowspec, and I2RS.

The above NSP functions are provided in a number of modules that can be used together or separately as shown in the following figure.

Figure 41: NSP functional modules



26698

The two main components of the NSP are:

- **Network Services Director (NSD)**


The NSD is a programmable and multi-vendor service provisioning tool that provides a single and simple API to the user and OSS. It implements a service model abstraction and adapts to each vendor-specific service model. It supports provisioning services such as E-Line, E-LAN, E-Tree, Layer 3 VPN, traffic steering, and service chaining.

- **Network Resource Controller (NRC)**

The NRC implements separate modules for computing and managing optimal paths for optical tunnels (NRC-T) and MPLS tunnels (NRC-P), and for computing optimal routing and placement of IP flows (NRC-F). In addition, a resource controller for inter-layer IP and optical path computation and more complex inter-domain MPLS path computation is provided as part of the Network Resource Controller Cross Domain (NRC-X).

The Network Resource Controller - Packet (NRC-P) implements the stateful PCE for packet networks. The following figure shows the NRC-P architecture and its main components.

Figure 42: NRC-P architecture



26697

The NRC-P has the following architecture:

- a single Virtual Machine (VM) handling the Java implementation of an MPLS path computation engine, a TE graph database, and an LSP database
- a plug-in adapter with the Nokia CPROTO interface, providing reliable, TCP-based message delivery between vSROS and Java-VM. The plug-in adapter implements a compact encoding/decoding (codec) function for the message content using Google ProtoBuf. Google ProtoBuf also provides for automatic C++ (vSROS side) and Java (Java-VM side) code generation to process the exchanged message content.
- a single VM running a vSROS image that handles the functions of topology discovery of multiple IGP instances and areas via IGP and NFM-P. For larger network domains, one VM running the vSROS image can be dedicated to a specific function.

The PCE module uses PCEP to communicate with its PCCs, and communicates with other PCEs to coordinate inter-domain path computation. Each router acting as a PCC initiates a PCEP session to the PCE in its domain.

When the user enables PCE control for one or more RSVP-TE LSPs, the PCE owns the path updating and periodic reoptimization of the LSPs. In this case, the PCE acts in an active stateful role. The PCE can also act in a passive stateful role for other LSPs on the router by discovering the LSPs and taking into account their resource consumption when computing the path for the LSPs it has control ownership of.

The following is a high-level description of the PCE and PCC capabilities:

- base PCEP implementation, as defined in RFC 5440
- active and passive stateful PCE LSP update, as defined in *draft-ietf-pce-stateful-pce*
- delegation of LSP control to the PCE

- synchronization of the LSP database with network elements for PCE-controlled LSPs and network element-controlled LSPs
- support for PCC-initiated LSPs, as defined in *draft-ietf-pce-stateful-pce*
- support for LSP path diversity across different LERs using extensions to the PCE path profile, as defined in *draft-alvarez-pce-path-profiles*
- support for LSP path bidirectionality constraints using extensions to the PCE path profile, as defined in *draft-alvarez-pce-path-profiles*

4.2 Base implementation of PCE

The base implementation of the PCE uses the PCEP extensions defined in RFC 5440.

The main functions of PCEP are:

- PCEP session establishment, maintenance, and closing
- path computation requests using the PCReq message
- path computation replies using the PCRep message
- notification messages (PCNtf) by which the PCEP speaker can inform its peer about events, such as path request cancellation by the PCC or path computation cancellation by the PCE
- error messages (PCErr) by which the PCEP speaker can inform its peer about errors related to processing requests, message objects, or TLVs

The following table lists the base PCEP TLVs, objects, and messages.

Table 40: Base PCEP TLVs, objects, and messages

TLV, object, or message	Contained in object	Contained in message
OPEN Object	N/A	OPEN, PCErr
Request Parameter (RP) Object	N/A	PCReq, PCRep, PCErr, PCNtf
NO-PATH Object	N/A	PCRep
END-POINTS Object	N/A	PCReq
BANDWIDTH Object	N/A	PCReq, PCRep, PCRpt ¹²
METRIC Object	N/A	PCReq, PCRep, PCRpt ¹²
Explicit Route Object (ERO)	N/A	PCRep
Reported Route Object (RRO)	N/A	PCRpt ¹²
LSPA Object	N/A	PCReq, PCRep, PCRpt ¹²

¹² Nokia proprietary

TLV, object, or message	Contained in object	Contained in message
Include Route Object (IRO)	N/A	PCReq, PCRep
SVEC Object	N/A	PCReq
NOTIFICATION Object	N/A	PCNtf
PCEP-ERROR Object	N/A	PCErr
LOAD-BALANCING Object	N/A	PCReq
CLOSE Object	N/A	CLOSE

The behavior and limitations of the implementation of the objects in the preceding table are as follows:

- The PCE treats all supported objects received in a PCReq message as mandatory, regardless of whether the P-flag in the object's common header is set (mandatory object) or not (optional object).
- The PCC implementation always sets the B-flag (B=1) in the metric object containing the hop metric value, which means that a bound value must be included in PCReq message. The PCE returns the computed value in the PCRep message with flags set identically to the PCReq message.
- The PCC implementation always sets flags B=0 and C=1 in the metric object for the IGP or TE metric values in the PCReq message. This means that the request is to optimize (minimize) the metric without providing a bound value. The PCE returns the computed value in the PCRep message with flags set identically to the PCReq message.
- The IRO and LOAD-BALANCING objects are not part of the NSP PCE feature. If the PCE receives a PCReq message with one or more of these objects, it ignores them regardless of the setting of the P-flag, and processes the path computations normally.
- The LSPA, metric, and bandwidth objects are also included in the PCRpt message. The inclusion of these objects in the PCRpt message is proprietary to Nokia.

The following features are not supported on the 7210 SAS:

- PCE discovery using IS-IS, as defined in RFC 5089, and OSPF, as defined in RFC 5088, along with corresponding extensions for discovering stateful PCE, as defined in *draft-sivabalan-pce-disco-stateful*
- security of the PCEP session using MD5 or TLS between PCEP peers
- PCEP synchronization optimization as defined in *draft-ietf-pce-stateful-sync-optimizations*
- support of end-to-end secondary backup paths for an LSP. PCE standards do not currently support an LSP container with multiple paths, and the PCE treats each request as a path with a unique PLSP-ID. It is up to the router to tie the two paths together to create 1:1 protection and to request path or SRLG diversity among them when it makes the request to the PCE.
- jitter, latency, and packet loss metrics support as defined in RFC 7471 and *draft-ietf-isis-te-metric-extensions*, and their use in the PCE metric object as defined in *draft-ietf-pce-pcep-service-aware*

4.3 PCEP session establishment and maintenance

PCEP operates over TCP using destination TCP port 4189. The PCC always initiates the connection. When the user configures the PCEP local address and the peer address on the PCC, the PCC initiates a TCP connection to the PCE. When a connection is established, the PCC and PCE exchange OPEN messages, which initializes the PCEP session and exchanges the session parameters to be negotiated.


The PCC always checks first to determine if the remote PCE address is reachable out-of-band through the management port. If the remote address is not reachable, the PCC tries to reach the remote PCE address in-band. When the session comes up out-of-band, the management IP address is always used; the local address configured by the user is ignored and is only used for an in-band session.

A keepalive mechanism is used as an acknowledgment of the acceptance of the session within the negotiated parameters. It is also used as a maintenance function to detect whether the PCEP peer is still alive.

The negotiated parameters include the **keepalive** timer and the **dead-timer**, and one or more PCEP capabilities such as support of stateful PCE and the LSP Path type.

The following figure shows the PCEP session initialization steps.

Figure 43: PCEP session initialization



If the session to the PCE times out, the router acting as a PCC keeps the last successfully programmed path provided by the PCE until the session to the PCE is reestablished. Any subsequent change to the LSP state is synchronized at the time the session is reestablished.

When a PCEP session to a peer times out or closes, the rate at which the PCEP speaker attempts to reestablish the session is subject to an exponential back-off mechanism.

4.4 PCEP parameters

The following PCEP parameters are user-configurable on the PCC:

- **keepalive timer**

A PCEP speaker must send a keepalive message if no other PCEP message is sent to the peer at the expiry of this timer. This timer is restarted every time a PCEP message is sent or the keepalive message is sent.

The keepalive mechanism is asymmetric, which allows each peer to use a different keepalive timer value.

The range of this timer is 1 to 255 seconds and the default value is 30 seconds.

- **dead timer**

This timer tracks the amount of time a PCEP speaker waits after the receipt of the last PCEP message before declaring its peer down.

The dead timer mechanism is asymmetric, which allows each PCEP speaker to propose a different dead timer value to its peer to detect session timeouts.

The range of this timer is 1 to 255 seconds and the default value is 120 seconds.

- **maximum rate of unknown messages**

When the rate of received unrecognized or unknown messages reaches the configured limit, the PCEP speaker closes the session to the peer.

The range of this message rate is 1 to 255 messages per minute and the default value is 10 messages per minute.

- **session reestablishment and state timeout**

If the PCEP session to the PCE goes down, all delegated PCC-initiated LSPs have their state maintained in the PCC and are not timed out. The PCC continues to try reestablishing the PCEP session. When the PCEP session is reestablished, the LSP database is synchronized with the PCE database, and any LSP that went down since the last time the PCEP session was up has its path updated by the PCE.

4.4.1 PCC configuration

The following PCC parameters can be modified while the PCEP session is operational:

- **report-path-constraints**
- **unknown-message-rate**

The following PCC parameters cannot be modified while the PCEP session is operational:

- **local-address**
- **keepalive**
- **dead-timer**
- **peer** (regardless of **shutdown** state)

4.4.2 LSP initiation

An LSP that is configured on the router is referred to as a PCC-initiated LSP. An LSP that is not configured on the router, but is instead created by the PCE at the request of an application or a service instantiation, is referred to as a PCE-initiated LSP.

The 7210 SAS supports three modes of operation for PCC-initiated LSPs, which are configurable on a per-LSP basis:

- **PCC-initiated and PCC-controlled**

When the path of the LSP is computed and updated by the router acting as a PCE Client (PCC), the LSP is referred to as PCC-initiated and PCC-controlled.

A PCC-initiated and PCC-controlled LSP has the following characteristics:

- The LSP can contain strict or loose hops, or a combination of both.
- CSPF is supported for RSVP-TE LSPs. Local path computation takes the form of hop-to-label translation for LSPs.
- LSPs can be reported to synchronize the LSP database of a stateful PCE server using the **pce-report** option. In this case, the PCE acts in passive stateful mode for this LSP. The LSP path cannot be updated by the PCE. The control of the LSP is maintained by the PCC.

- **PCC-initiated and PCE-computed**

When the path of the LSP is computed by the PCE at the request of the PCC, it is referred to as PCC-initiated and PCE-computed.

A PCC-initiated and PCE-computed LSP has the following characteristics:

- The **pce-computation** option must be enabled for the LSP so that the PCE can perform path computation at the request of the PCC only. The PCC retains control.
- LSPs can be reported to synchronize the LSP database of a stateful PCE server using the **pce-report** option. In this case, the PCE acts in passive stateful mode for this LSP.

- **PCC-initiated and PCE-controlled**

When the path of the LSP is updated by the PCE following a delegation from the PCC, it is referred to as PCC-initiated and PCE-controlled.

A PCC-initiated and PCE-controlled LSP has the following characteristics:

- The **pce-control** option must be enabled for the LSP so that the PCE can perform path updates following a network event without an explicit request from the PCC. The PCC delegates full control.
- The **pce-report** option must be enabled for LSPs that cannot be delegated to the PCE. The PCE acts in active stateful mode for this LSP.

4.4.3 PCC-initiated and PCE-computed or PCE-controlled LSPs

The following is the procedure for configuring and programming a PCC-initiated LSP when control is delegated to the PCE.

1. The LSP configuration is created on the PE router via CLI or via the OSS/NSP NFM-P.
The configuration dictates which PCE control mode is wanted: active (**pce-control** and **pce-report** options enabled) or passive (**pce-computation** enabled and **pce-control** disabled).
2. PCC assigns a unique PLSP-ID to the LSP. The PLSP-ID uniquely identifies the LSP on a PCEP session and must remain constant during its lifetime. PCC on the router must keep track of the association of the PLSP-ID to the Tunnel-ID and Path-ID, and use the latter to communicate with MPLS about a specific path of the LSP. PCC also uses the SRP-ID to correlate PCRpt messages for each new path of the LSP.

3. The PE router does not validate the entered path. However, in the 7210 SAS, the PCE supports the computation of a path for an LSP with empty-hops in its path definition. While PCC will include the IRO objects in the PCReq message to PCE, the PCE will ignore them and compute the path with the other constraints except the IRO.
4. The PE router sends a PCReq message to the PCE to request a path for the LSP, and includes the LSP parameters in the METRIC object, the LSPA object, and the BANDWIDTH object. The PE router also includes the LSP object with the assigned PLSP-ID. At this point, the PCC does not delegate the control of the LSP to the PCE.
5. The PCE computes a new path, reserves the bandwidth, and returns the path in a PCRep message with the computed ERO in the ERO object. It also includes the LSP object with the unique PLSP-ID, the METRIC object with any computed metric value, and the BANDWIDTH object.



Note:

To enable the PCE to use the SRLG path diversity and admin-group constraints in the path computation, the user must configure the SRLG and admin-group membership against the MPLS interface and enable the **traffic-engineering** option in IGP. This causes IGP to flood the link SRLG and admin-group membership in its participating area, and for the PCE to learn it in its TE database.

6. The PE router updates the CPM and the datapath with the new path.
Up to this step, the PCC and PCE are using passive stateful PCE procedures. The next steps will synchronize the LSP database of the PCC and PCE for both PCE-computed and PCE-controlled LSPs. They will also initiate the active PCE stateful procedures for the PCE-controlled LSP only.
7. The PE router sends a PCRpt message to update the PCE with an Up state, and also sends the RRO as confirmation. It now includes the LSP object with the unique PLSP-ID. For a PCE-controlled LSP, the PE router also sets the delegation control flag to delegate control to the PCE. The state of the LSP is now synchronized between the router and the PCE.
8. Following a network event or a reoptimization, the PCE computes a new path for a PCE-controlled LSP and returns it in a PCUpd message with the new ERO. It will include the LSP object with the same unique PLSP-ID assigned by the PCC, as well as the Stateful Request Parameter (SRP) object with a unique SRP-ID-number to track error and state messages specific to this new path.
9. The PE router updates the CPM and the datapath with the new path.
10. The PE router sends a PCRpt message to inform the PCE that the older path is deleted. It includes the unique PLSP-ID value in the LSP object and the R (Remove) bit set.
11. The PE router sends a new PCRpt message to update PCE with an Up state, and also sends the RRO to confirm the new path. The state of the LSP is now synchronized between the router and the PCE.
12. If PCE owns the delegation of the LSP and is making a path update, MPLS will initiate the LSP and update the operational value of the changed parameters while the configured administrative values will not change. Both the administrative and operational values are shown in the details of the LSP path in MPLS.
13. If the user makes any configuration change to the PCE-computed or PCE-controlled LSP, MPLS requests that the PCC first revoke delegation in a PCRpt message (PCE-controlled only), and then MPLS and PCC follow the above steps to convey the changed constraint to PCE which will result in the programming of a new path into the datapath, the synchronization of the PCC and PCE LSP databases, and the return of delegation to PCE.

The preceding procedure is followed when the user performs a **no shutdown** command on a PCE-controlled or PCE-computed LSP. The starting point is an LSP which is administratively down with no active path. For an LSP with an active path, the following items may apply:

1. If the user has enabled the **pce-computation** option on a PCC-controlled LSP with an active path, no action is performed until the next time the router needs a path for the LSP following a network event of a LSP parameter change. At that point, the prior procedure is followed.
2. If the user has enabled the **pce-control** option on a PCC-controlled or PCE-computed LSP with an active path, the PCC will issue a PCRpt message to the PCE with an Up state, as well as the RRO of the active path. It will set the delegation control flag to delegate control to the PCE. The PCE will keep the active path of the LSP and make no updates to it until the next network event or reoptimization. At that point, the prior procedure is followed.

4.5 PCEP support for RSVP-TE LSPs

This section describes the support of PCC-initiated RSVP-TE LSPs. PCEP support of an RSVP-TE LSP is described in [LSP initiation](#) with the following differences:

- each primary and secondary path is assigned its own unique path LSP-ID (PLSP-ID)
- the PCC indicates to the PCE the state of each path (either up or down) and which path is currently active and carrying traffic (active state)

4.5.1 RSVP-TE LSP configuration for a PCC router

The following MPLS-level and LSP-level CLI commands are used to configure RSVP-TE LSPs in a router acting as a PCEP Client (PCC):

- ```
config>router>mpls>
 pce-report rsvp-te {enable | disable}
```
- ```
config>router>mpls>lsp>
    path-profile profile-id [path-group group-id]
    pce-computation
    pce-control
    pce-report {enable | disable | inherit}
```

The **cspf** option must be enabled on the LSP before the **pce-computation** or **pce-control** options can be enabled. An attempt to disable the **cspf** option on an RSVP-TE LSP that has the **pce-computation** or **pce-control** options enabled will be rejected.

If the LSP has disabled PCE reporting, either because of inheritance from the MPLS-level configuration or because of LSP-level configuration, enabling the **pce-control** option for the LSP has no effect. To help troubleshoot this situation, the output of the **show** commands for the LSP displays the operational values of both the **pce-report** and **pce-control** options.



Note:

The PCE function implemented in the NSP and referred to as the NRC-P, supports only Shared Explicit (SE) style bandwidth management for RSVP-TE LSPs. The PCEP does not support the ability of the PCC to convey this value to the PCE. Therefore, whether the LSP configuration

option **rsvp-resv-style** is set to **se** or **ff**, the PCE will always use the SE style in the CSPF computation of the path for a PCE-computed or PCE-controlled RSVP-TE LSP.

A manual bypass LSP does not support any of the PCE-related commands. Reporting a bypass LSP to the PCE is not required because the bypass LSP does not book bandwidth.

All other MPLS, LSP, and path-level commands are supported, with the exception of the following commands:

- **least-fill**
- **srlg** (on secondary standby path)

For more information about RSVP-TE PCC instantiation modes, see [LSP initiation](#).

4.5.2 Behavior of the LSP path update

When the **pce-control** option is enabled, the PCC delegates control of the RSVP-TE LSP to the PCE.

The NRC-P sends a path update using the PCUpd message in the following cases:

- a failure event that impacts a link or a node in the path of a PCE-controlled LSP
The operation is performed by the PCC as a Make-Before-Break (MBB) if the LSP remained in the up state because of protection provided by FRR or a secondary path. If the LSP went down, the update brings it into the up state. A PCRpt message is sent by the PCC for each change to the state of the LSP during this process. See [Behavior of LSP MBB](#) for more information.
- a topology change that impacts a link in the path of a PCE-controlled LSP
This topology change can be a change to the IGP metric, the TE metric, admin-group, or SRLG membership of an interface. This update is performed as an MBB by the PCC.
- the user has performed a manual resignal of a PCE-controlled RSVP-TE LSP path from the NRC-P
This update is performed as an MBB by the PCC.
- the user has performed a Global Concurrent Optimization (GCO) on a set of PCE-controlled RSVP-TE LSPs from the NRC-P

This update is performed as an MBB by the PCC.

The procedures for the path update are described in [LSP initiation](#). However, for an RSVP-TE LSP, the PCUpd message from the PCE contains the interface IP address or system IP address in the computed ERO. The PCC signals the path using the ERO returned by the PCE and, if successful, programs the datapath, then sends the PCRpt message with the resulting RRO and hop labels provided by RSVP-TE signaling.

If the signaling of the ERO fails, the ingress LER returns a PCErr message to the PCE with the LSP Error code field of the LSP-ERROR-CODE TLV set to a value of 8 (RSVP signaling error).

If the **no adaptive** option is set for the RSVP-TE LSP, the ingress LER cannot perform an MBB for the LSP. A PCUpd message received from the PCE is then failed by the ingress LER, which returns a PCErr message to the PCE with the LSP Error code field of the LSP-ERROR-CODE TLV set to a value of 8 (RSVP signaling error).

4.5.2.1 Path update with empty ERO

When the NRC-P reoptimizes the path of a PCE-controlled RSVP-TE LSP, it is possible that a path that satisfies the constraints of the LSP no longer exists. In this case, the NRC-P sends a PCUpd message with an empty ERO, which forces the PCC to bring down the path of the RSVP-TE LSP.

The NRC-P sends a PCUpd message with an empty ERO if any of the following cases are true:

- the requested bandwidth is the same as the current bandwidth, which avoids bringing down the path because of a resignal during an MBB transition
- local protection is not currently in use, which avoids bringing down a path that activated an FRR backup path. The LSP can remain on the FRR backup path until a new primary path can be found by the NRC-P.
- the links of the current path are all operationally up, which allows the NRC-P to ensure that the RSVP control plane will report the path down when a link is down and not prematurely bring the path down with an empty ERO

4.5.3 Behavior of LSP MBB

In addition to the MBB support when the PCC receives a path update, as described in [Behavior of the LSP path update](#), an RSVP-TE LSP supports the MBB procedure for any parameter configuration change, including the PCEP-related commands when they result in a change to the path of the LSP.

If the user adds or modifies the **path-profile** command for an RSVP-TE LSP, a configuration change MBB is only performed if the **pce-computation**, **pce-report**, or **pce-control** options are enabled on the LSP. Otherwise, no action occurs. When **pce-computation**, **pce-report**, or **pce-control** are enabled on the LSP, the path update MBB (**tools>perform>router>mpls>update-path**) fails, resulting in no operation.

MBB is also supported for the manual resignal MBB type.

If the LSP goes into an MBB state at the ingress LER, the behavior is dependent on the operating mode of the LSP.

4.5.3.1 PCC-controlled LSPs

All MBB types are supported for PCC-controlled LSPs. The LSP MBB procedures for a PCC-controlled LSP (**pce-computation** and **pce-control** disabled) are as follows:

1. MPLS submits a path request, including the updated path constraints, to the local CSPF.
2. If the local CSPF returns a path, the PCC signals the LSP with the RSVP control plane and moves traffic to the new MBB path. If **pce-report** is enabled for this LSP, the PCC sends a PCRpt message with the delegation bit clear to retain control and containing the RRO and LSP objects, with the LSP-IDENTIFIERS TLV containing the LSP-ID of the new MBB path. The message includes the metric, LSPA, and bandwidth objects where the P-flag is clear, which indicates the operational values of these parameters. Unless the user disables the **report-path-constraints** option under the **pcc** context, the PCC also includes a second set of metric, LSPA, and bandwidth objects with the P-flag set to convey to the PCE the constraints of the path.
3. If the CSPF returns no path or the RSVP-TE signaling of the returned path fails, MPLS puts the LSP into retry mode and sends a request to the local CSPF every *retry-timer* seconds and up to the value of *retry-count*.

4. When **pce-report** is enabled for the LSP and the FRR global revertive MBB is triggered following a bypass LSP activation by a PLR in the network, the PCC issues an updated PCRpt message with the new RRO reflecting the PLR and RRO hops. The PCE releases the bandwidth on the links that are no longer used by the LSP path.

4.5.3.2 PCE-computed LSPs

All MBB types are supported for PCE-computed LSPs. The LSP MBB procedures for a PCE-computed LSP (**pce-computation** enabled and **pce-control** disabled) are as follows:

1. The PCC issues a PCReq for the same PLSP-ID and includes the updated constraints in the metric, LSPA, and bandwidth objects.
 - If the PCE successfully finds a path, it replies with a PCRep message with the ERO.
 - If the PCE does not find a path, it replies with a PCRep message containing the No-Path object.
2. If the PCE returns a path, the PCC signals the LSP with the RSVP control plane and moves traffic to the new MBB path. If **pce-report** is enabled for this LSP, the PCC sends a PCRpt message with the delegation D-bit clear to retain control and containing the RRO and LSP objects, with the LSP-IDENTIFIERS TLV containing the LSP-ID of the new MBB path. The message includes the metric, LSPA, and bandwidth objects where the P-flag is clear, which indicates the operational values of these parameters. Unless the user disables the **report-path-constraints** option under the **pcc** context, the PCC also includes a second set of metric, LSPA, and bandwidth objects with the P-flag set to convey to the PCE the constraints of the path.
3. If the PCE returns no path or the RSVP-TE signaling of the returned path fails, MPLS puts the LSP into retry mode and sends a request to PCE every *retry-timer* seconds and up to the value of *retry-count*.
4. When the **pce-report** is enabled for the LSP and the FRR global revertive MBB is triggered following a bypass LSP activation by a PLR in the network, the PCC issues an updated PCRpt message with the new RRO reflecting the PLR and RRO hops. The PCE releases the bandwidth on the links that are no longer used by the LSP path.
5. If the user changes the RSVP-TE LSP configuration from **pce-computation** to **no pce-computation**, MBB procedures are not supported. In this case, the LSP path is torn down and is put into retry mode to compute a new path from the local CSPF on the router to signal the LSP.

4.5.3.3 PCE-controlled LSPs

The LSP MBB procedures for a PCE-controlled LSP (**pce-control** enabled) are as follows.



Note:

Items 1 through 5 of the following procedure apply to the config change, and manual resignal MBB types. The delayed retry MBB type used with the SRLG on secondary standby LSP feature is not supported with a PCE-controlled LSP. See [Behavior of secondary LSP Paths](#) for information about the SRLG on secondary standby LSP feature.

1. The PCC temporarily removes delegation by sending a PCRpt message for the corresponding path LSP-ID (PLSP-ID) with the delegation D-bit clear.
2. For an LSP with **pce-computation** disabled, MPLS submits a path request to the local CSPF, which includes the updated path constraints.

3. For an LSP with **pce-computation** enabled, the PCC issues a PCReq for the same PLSP-ID and includes the updated constraints in the metric, LSPA, or bandwidth objects.
 - If the PCE successfully finds a path, it replies with a PCRep message with the ERO.
 - If the PCE does not find a path, it replies with a PCRep message containing the No-Path object.
4. If the local CSPF or the PCE returns a path, the PCC performs the following actions:
 - The PCC signals the LSP with the RSVP control plane and moves traffic to the new MBB path. It then sends a PCRpt message with the delegation D-bit set to return delegation and containing the RRO and LSP objects, with the LSP-IDENTIFIERS TLV containing the LSP-ID of the new MBB path. The message includes the metric, LSPA, and bandwidth objects where the P-flag is clear, which indicates the operational values of these parameters. Unless the user disabled the **report-path-constraints** option under the **pcc** context, the PCC also includes a second set of metric, LSPA, or bandwidth objects with the P-flag set to convey to the PCE the constraints of the path.
 - The PCC sends a PathTear message to delete the state of the older path in the network. The PCC then sends a PCRpt message to the PCE with the older path LSP (PLSP-ID) and the remove R-bit set to also have the PCE remove the state of that LSP from its database.
5. If the local CSPF or the PCE returns no path or the RSVP-TE signaling of the returned path fails, the router makes no further requests. That is, there is no retry for the MBB.
 - The PCC sends a PCErr message to the PCE with the LSP Error code field of the LSP-ERROR-CODE TLV set to a value of 8 (RSVP signaling error) if the MBB failed because of a RSVP-TE signaling error.
 - The PCC sends a PCRpt message with the delegation D-bit set to return delegation and containing the RRO and LSP objects, with the LSP-IDENTIFIERS TLV containing the LSP-ID of the currently active path. The message includes the metric, LSPA, and bandwidth objects where the P-flag is clear to indicate the operational values of these parameters. Unless the user disabled the **report-path-constraints** option under the **pcc** context, the PCC also includes a second set of metric, LSPA, and bandwidth objects with the P-flag set to convey to the PCE the constraints of the path.
6. The ingress LER takes no action in the case of a network event triggered MBB, such as FRR global revertive or TE graceful shutdown.
 - The ingress PE keeps the information as required and sets the state of MBB to one of the FRR global revertive or TE graceful shutdown MBB values but does not perform the MBB action.
 - The NRC-P computes a new path for the global revertive MBB because of a failure event. This computation uses the PCUpd message to update the path using the MBB procedure described in [Behavior of the LSP path update](#). The activation of a bypass LSP by a point of local repair (PLR) in the network causes the PCC to issue an updated PCRpt message with the new RRO reflecting the PLR and RRO hops. The PCE will release the bandwidth on the links that are no longer used by the LSP path.
 - The NRC-P computes a new path for the TE graceful shutdown MBB if the RSVP-TE is using the TE metric, because the TE metric of the link in TE graceful shutdown is set to infinity. This computation uses the PCUpd message to update the path using the MBB procedure described in [Behavior of the LSP path update](#).
 - The NRC-P does not act on the TE graceful shutdown MBB if the RSVP-TE is using the IGP metric; however, the user can perform a manual resignal of the LSP path from the NRC-P to force a new path computation, which accounts for the newly available bandwidth on the link that caused the MBB event. This computation uses the PCUpd message to update the path using the MBB procedure described in [Behavior of the LSP path update](#).

- The user can perform a manual resignal of the LSP path from the ingress LER, which forces an MBB for the path as per the remove-delegation/MBB/return-delegation procedures described in this section.
 - If the user performs **no pce-control** while the LSP still has the state for any of the network event triggered MBBs, the MBB is performed immediately by the PCC as described in the procedures in [PCE-computed LSPs](#) for a PCE-computed LSP and as described in the procedures in [PCC-controlled LSPs](#) for a PCC-controlled LSP.
7. The timer-based manual resignal MBB behaves like the TE graceful shutdown MBB. The user can perform a manual resignal of the LSP path from the ingress LER or from the PCE.
 8. The path update MBB (**tools>perform>router>mpls>update-path**) fails, which results in no operation. This is true in all cases when the RSVP-TE LSP enables the **pce-report** option.

4.5.4 Behavior of secondary LSP Paths

Each of the primary, secondary standby, and secondary non-standby paths of the same LSP must use a separate path LSP-ID (PLSP-ID). The PCE function of the NSP, the NRC-P, checks the LSP-IDENTIFIERS TLV in the LSP object and can identify which PLSP-IDs are associated with the same LSP or the same RSVP-TE session. The parameters are the IPv4 Tunnel Sender Address, the Tunnel ID, the Extended Tunnel ID, and the IPv4 Tunnel Endpoint Address. This approach allows the use of all the PCEP procedures for all three types of LSP paths.

The PCC indicates to the PCE the following states for the path in the LSP object: down, up (signaled but not carrying traffic), or active (signaled and carrying traffic).

The PCE tracks active paths and displays them in the NSP GUI. It also provides only the tunnel ID of an active PLSP-ID to a destination prefix when a request is made by a service or a steering application.

The PCE recomputes the paths of all PLSP-IDs that are affected by a network event. The user can select each path separately on the NSP GUI and trigger a manual resignal of one or more paths of the RSVP-TE LSP.



Note:

Enabling the **srlg** option on a secondary standby path results in no operation. The NRC-P supports link and SRLG disjointedness using the PCE path profile. The user can apply the PCE path profile to the primary and secondary paths of the same LSP. See [PCE path profile support](#) for more information.

4.5.5 PCE path profile support

The PCE path profile ID and path group ID are configured at the LSP level (**config>router>mpls>lsp>path-profile**).

The NRC-P can enforce path disjointedness and bidirectionality among a pair of forward and a pair of reverse LSP paths. Both pairs of LSP paths must use a unique path group ID along with the same path profile ID.

If the user wants to apply path disjointedness and path bidirectionality constraints to RSVP-TE LSP paths, it is important to follow the following guidelines. The user can configure the following sets of LSP paths:

- a set consisting of a pair of forward RSVP-TE LSPs and a pair of reverse RSVP-TE LSPs, each with a single primary or secondary path. The pair of forward LSPs can originate and terminate on different

routers. The pair of reverse LSPs must mirror the forward pair. In this case, the path profile ID and the path group ID configured for each LSP must match. Because each LSP has a single path, the bidirectionality constraint applies automatically to the forward and reverse LSPs, which share the same originating node and the same terminating routers.

- a pair consisting of a forward RSVP-TE LSP and a reverse RSVP-TE LSP, each with a primary path and a single secondary path, or each with two secondary paths. Because the two paths of each LSP inherit the same LSP level path profile ID and path group ID configuration, the NRC-P path computation algorithm cannot guarantee that the primary paths in both directions meet the bidirectionality constraint. That is, it is possible that the primary path for the forward LSP shares the same links as the secondary path of the reverse LSP and the other way around.

4.6 LSP path diversity and bidirectionality constraints

The PCE path profile defined in *draft-alvarez-pce-path-profiles* is used to request path diversity or a disjoint for two or more LSPs originating on the same or different PE routers. It is also used to request that paths of two unidirectional LSPs between the same two routers use the same TE links. This is referred to as the bidirectionality constraint.

Path profiles are defined directly on the NRC-P Policy Manager with a number of LSP path constraints, which are metrics with upper bounds specified, and with an objective, which are metrics optimized with no bounds specified. The NRC-P Policy Manager allows the following PCE constraints to be configured within each PCE path profile:

- path diversity, node-disjoint, link-disjoint
- path bidirectionality, symmetric reverse route preferred, symmetric reverse route required
- maximum path IGP metric (cost)
- maximum path TE metric
- maximum hop count

The user can also specify the PCE objective used to optimize the path of the LSP in the PCE path profile, one of:

- IGP metric (cost)
- TE metric
- hops (span)

The CSPF algorithm will optimize the objective. If a constraint is provided for the same metric, the CSPF algorithm ensures that the selected path achieves a lower or equal value to the bound value specified in the constraint.

For hop-count metrics, if a constraint is sent in a metric object and is also specified in a PCE profile referenced by the LSP, the constraint in the metric object is used.

For IGP and TE metrics, if an objective is sent in a metric object and is also specified in a PCE profile referenced by the LSP, the objective in the path profile is used.

The constraints in the bandwidth object and the LSPA object, specifically the include and exclude admin-group constraints and setup and hold priorities, are not supported in the PCE profile.

To indicate the path diversity and bidirectionality constraints to the PCE, the user must configure the profile ID and path group ID of the PCE path to which the LSP belongs. The path group ID does not need to be

defined in the PCE as part of the path profile configuration and identifies implicitly the set of paths that must have the path diversity constraint applied.

The user can only associate a single path group ID with a specific PCE path profile ID for an LSP. However, the same path group ID can be associated with multiple PCE profile IDs for the same LSP.

The path profiles are inferred using the path ID in the path request by the PCC. When the PE router acting as a PCC wants to request path diversity from a set of other LSPs belonging to a path group ID value, it adds a new PATH-PROFILE object in the PCReq message. The object contains the path profile ID and the path group ID as an extended ID field. In other words, the diversity metric is carried in an opaque way from the PCC to the PCE.

The bidirectionality constraint operates the same way as the diversity constraint. The user can configure a PCE profile with both the path diversity and bidirectionality constraints. The PCE will check if there is an LSP in the reverse direction that belongs to the same path group ID as an originating LSP it is computing the path for, and will enforce the constraint.

To ensure that the PCE is aware of the path diversity and bidirectionality constraints for an LSP that is delegated but for which there is no prior state in the NRC-P LSP database, the PATH-PROFILE object is included in the PCRpt message with the P-flag set in the common header to indicate that the object must be processed. This is proprietary to Nokia.

The following table lists the new objects and TLVs introduced in the PCE path profile.

Table 41: PCEP path profile extension objects and TLVs

TLV, object, or message	Contained in object	Contained in message
PATH-PROFILE-CAPABILITY TLV	OPEN	OPEN
PATH-PROFILE Object	N/A	PCReq, PCRpt ¹³

A PATH-PROFILE object can contain multiple TLVs containing each profile ID and extend ID, and should be processed properly. If multiple PATH-PROFILE objects are received, the first object is interpreted and the others are ignored. The PCC and the PCE support all PCEP capability TLVs defined in this document and will always advertise them. If the OPEN object received from a PCEP speaker does not contain one or more of the capabilities, the PCE or PCC will not use them during that PCEP session.

4.7 PCEP configuration command reference



Note:
PCEP commands are only supported on the 7210 SAS-Mxp, 7210 SAS-R6, and 7210 SAS-R12.

4.7.1 Command hierarchies

- [PCEP commands](#)
- [Show commands](#)

¹³ Nokia proprietary

- [Tools commands](#)

4.7.1.1 PCEP commands

```
config
- router
  - [no] pcep
    - [no] pcc
      - dead-timer seconds
      - no dead-timer
      - keepalive seconds
      - no keepalive
      - local-address ip-address
      - no local-address
      - [no] peer ip-address
        - [no] shutdown
      - [no] report-path-constraints
      - [no] shutdown
      - unknown-message-rate msg/min
      - no unknown-message-rate
```

4.7.1.2 Show commands

```
show
- router
  - pcep
    - pcc
      - detail
      - lsp-db [lsp-type lsp_type] [delegated-pce ip-address]
      - lsp-db [lsp-type lsp_type] from ip-address [delegated-pce ip-address]
      - lsp-db [lsp-type lsp_type] lsp lsp-name [delegated-pce ip-address]
      - lsp-db [lsp-type lsp_type] to ip-address [tunnel-id [tunnel-id]]
      - lsp-db [lsp-type lsp_type] tunnel-id [tunnel-id]
      - path-request [lsp-type {rsvp-p2p}] [dest ip-address] [detail]
      - peer [ip-address] [detail]
      - status
```

4.7.1.3 Tools commands

```
tools
- dump
  - router
    - pcep
      - pcc lsp [plsp-id plsp-id]
      - pcc lsp lsp-type lsp-type [tunnel-id tunnel-id]
```

4.7.2 Command descriptions

- [PCEP commands](#)
- [Show commands](#)
- [Tools commands](#)

4.7.2.1 PCEP commands

pcep

Syntax

[no] pcep

Context

config>router

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

This command enables the Path Computation Element Communication Protocol (PCEP) and enters the context to configure PCEP parameters.

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

pcc

Syntax

[no] pcc

Context

config>router>pcep

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

Commands in this context configure PCC parameters.

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

dead-timer

Syntax

dead-timer *seconds*

no dead-timer

Context

config>router>pcep>pcc

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

This command configures the PCEP session dead timer value, which is the amount of time a PCEP speaker will wait after the receipt of the last PCEP message before declaring its peer down.

The dead timer mechanism is asymmetric, which means that each PCEP speaker can propose a different dead timer value to its peer to use to detect session timeout.

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

Default

dead-timer 120

Parameters

seconds

Specifies the dead timer value, in seconds.

Values 1 to 255

keepalive

Syntax

keepalive *seconds*

no keepalive

Context

config>router>pcep>pcc

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

This command configures the PCEP session keepalive value. A PCEP speaker must send a keepalive message if no other PCEP message is sent to the peer at the expiry of this timer. This timer is restarted every time a PCEP message or keepalive message is sent.

The keepalive mechanism is asymmetric, which means that each peer can use a different keepalive timer value at its end.

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

Default

keepalive 30

Parameters

seconds

Specifies the keepalive value, in seconds.

Values 1 to 255

local-address

Syntax

local-address *ip-address*

no local-address

Context

config>router>pcep>pcc

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

This command configures the local address of the PCEP speaker.

The PCEP operates over TCP using destination TCP port 4189. The PCE client (PCC) always initiates the connection. When the user configures the PCEP local address and the peer address on the PCC, the PCC initiates a TCP connection to the PCE. When the connection is established, the PCC and PCE exchange OPEN messages, which initializes the PCEP session and exchanges the session parameters to be negotiated.

The PCC always checks first to determine if the remote PCE address is reachable out-of-band through the management port. If the remote address is not reachable, the PCC checks if the remote PCE address is reachable in-band. If the session comes up out-of-band, the system IP address is always used. The local address configured by the user is only used for in-band sessions and is otherwise ignored.

The **no** form of this command removes the configured local address of the PCEP speaker.

Parameters

ip-address

Specifies the IP address of the PCEP speaker to be used for in-band sessions.

peer

Syntax

[no] peer *ip-address*

Context

config>router>pcep>pcc

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

This command configures the IP address of a peer PCEP speaker. The address is used as the destination address in the PCEP session messages to a PCEP peer.

The **no** form of this command removes the specified peer PCEP speaker.

Parameters

ip-address

Specifies the IP address of the PCEP peer to be used as the destination address in the PCEP session.

Values a.b.c.d

report-path-constraints

Syntax

[no] report-path-constraints

Context

config>router>pcep>pcc

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

This command enables the inclusion of LSP path constraints in the PCE report messages sent from the PCC to a PCE.

For the PCE to know about the original constraints for an LSP that is delegated but for which there is no prior state in its LSP database (for example, if no PCReq message was sent for the same PLSP-ID), the following proprietary behavior is observed.

- The PCC appends a duplicate of each of the LSPA, metric, and bandwidth objects in the PCRpt message. The only difference between two objects of the same type is that the P-flag is set in the common header of the duplicate object to indicate that it is a mandatory object for processing by the PCE.
- The value of the metric or bandwidth in the duplicate object contains the original constraint value, while the first object contains the operational value. This is applicable to hop metrics in the metric and bandwidth objects only. The 7210 SAS PCC does not support configuring a boundary on the path computation IGP or TE metrics.

- The path computation on the PCE must use the first set of objects when updating a path if the PCRpt message contained a single set. If the PCRpt message contained a duplicate set, PCE path computation must use the constraints in the duplicate set.

The **no** form of this command disables the preceding behavior in case of interoperability issues with third-party PCE implementations.

Default

report-path-constraints

shutdown

Syntax

[no] shutdown

Context

config>router>pcep>pcc

config>router>pcep>pcc>peer

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

This command administratively disables the PCC process.

The following PCC parameters can be modified without shutting down the PCEP session:

- **report-path-constraints**
- **unknown-message-rate**

The following PCC parameters can only be modified when the PCEP session is shut down:

- **local-address**
- **keepalive**
- **dead-timer**
- **peer**

The **no** form of this command administratively enables the PCC process.

Default

shutdown

Special Cases

PCEP Protocol Handling for 7210 SAS-Mxp

When the **no shutdown** command is issued in the **configure>router>pcep>pcc** context, resources are allocated to enable processing of the protocol by the node. When you issue the **configure>router>pcep>pcc>shutdown** command, the resources are deallocated.

unknown-message-rate

Syntax

unknown-message-rate *msg/min*

no unknown-message-rate

Context

config>router>pcep>pcc

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

This command configures the maximum rate of unknown messages that can be received during a PCEP session.

When the rate of received unrecognized or unknown messages reaches the configured limit, the PCEP speaker closes the session to the peer.

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

Default

unknown-message-rate 10

Parameters

msg/min

Specifies the rate of unknown messages, in messages per minute.

Values 1 to 255

4.7.2.2 Show commands



Note:

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

detail

Syntax

detail

Context

show>router>pcep>pcc

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

This command displays detailed information about PCEP PCC.

Output

The following output is an example of PCEP PCC more information, and [Table 42: Output fields: PCEP PCC](#) describes the output fields.

Sample output

```
*A:ZBG>show>router>pcep# pcc detail
=====
Path Computation Element Protocol (PCEP) Path Computation Client (PCC) Info
=====
Admin Status           : Down           Oper Status           : Down
Unknown Msg Limit      : 10 msg/min
Keepalive Interval     : 50 seconds   DeadTimer Interval    : 150 seconds
Capabilities List       : stateful-delegate stateful-pce rsvp-path
Address                : 10.10.10.10
Report Path Constraints: True
Open Wait Timer        : 60 seconds   Keep Wait Timer       : 60 seconds
Sync Timer             : 60 seconds   Request Timer        : 120 seconds
Connection Timer       : 60 seconds   Allow Negotiations    : False
Max Sessions           : 1           Max Unknown Req       : 1000
=====
*A:ZBG>show>router>pcep#
```

Table 42: Output fields: PCEP PCC

Label	Description
Admin Status	Displays the administrative status of the PCC
Oper Status	Displays the operational status of the PCC
Unknown Msg Limit	Displays the maximum rate of unknown messages that can be received on a PCEP session
Keepalive Interval	Displays the specified keepalive interval for the PCEP session
DeadTimer Interval	Displays the specified dead time interval for the PCEP session
Capabilities List	Displays the capabilities list for the PCEP session
Address	Displays the local IP address of the PCEP speaker
Report Path Constraints	Displays whether to include LSP path constraints in the PCE report messages sent from the PCC to a PCE
Open Wait Timer	Displays the value of the open wait timer for the PCEP session

Label	Description
Keep Wait Timer	Displays the value of the keep wait timer for the PCEP session
Sync Timer	Displays the value of the synchronization timer for the PCEP session
Request Timer	Displays the value of the request timer for the PCEP session
Connection Timer	Displays the value of the keep wait timer for the PCEP session
Allow Negotiations	Displays where negotiations between PCEP PCC and PCE are allowed
Max Sessions	Displays the maximum number of PCEP sessions on the router
Max Unknown Req	Displays the maximum number of unknown requests for PCEP sessions on the router

lsp-db

Syntax

```
lsp-db [lsp-type lsp_type] [delegated-pce ip-address]
lsp-db [lsp-type lsp_type] from ip-address [delegated-pce ip-address]
lsp-db [lsp-type lsp_type] lsp lsp-name [delegated-pce ip-address]
lsp-db [lsp-type lsp_type] to ip-address [tunnel-id [tunnel-id]]
lsp-db [lsp-type lsp_type] tunnel-id [tunnel-id]
```

Context

```
show>router>pcep>pcc
```

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

This command displays PCEP PCC LSP information.

Parameters

lsp_type

Specifies the type of LSP to display. The only available option is RSVP-TE point-to-point LSPs (rsvp-p2p).

tunnel-id

Specifies a tunnel ID.

Values 1 to 65535

ip-address

Specifies an IPv4 address.

Values a.b.c.d

Output

The following output is an example of PCEP PCC LSP information, and [Table 43: Output fields: LSP](#) describes the output fields.

Sample output

```
*A:ZBG# show router pcep pcc lsp-db
=====
PCEP Path Computation Client (PCC) LSP Update Info
=====
PCEP-specific LSP ID: 1
LSP ID           : 21504           LSP Type          : rsvp-p2p
Tunnel ID        : 1              Extended Tunnel Id   : 10.20.1.3
LSP Name         : test_lsp::fully_loose
Source Address   : 10.20.1.3       Destination Address : 10.20.1.1
LSP Delegated    : True            Delegate PCE Address: 138.120.210.36
Oper Status      : active
-----
PCEP-specific LSP ID: 2
LSP ID           : 21510           LSP Type          : rsvp-p2p
Tunnel ID        : 1              Extended Tunnel Id   : 10.20.1.3
LSP Name         : test_lsp::stdby_fully_loose_2
Source Address   : 10.20.1.3       Destination Address : 10.20.1.1
LSP Delegated    : True            Delegate PCE Address: 10.120.210.36
Oper Status      : up
=====
*A:ZBG#
```

Table 43: Output fields: LSP

Label	Description
PCEP-specific LSP ID	Displays the PCEP-specific LSP identifier
LSP ID	Displays the LSP identifier
Tunnel ID	Displays the tunnel identifier for the LSP
LSP Name	Displays the configured LSP name
Source Address	Displays the source IP address of the LSP
LSP Delegated	Displays the delegation status of the LSP
Oper Status	Displays the operational status of the LSP

Label	Description
LSP Type	Displays the type of the LSP
Extended Tunnel ID	Displays the expanded tunnel identifier for the LSP
Destination Address	Displays the destination IP address of the LSP
Delegate PCE Address	Displays the IP address of the delegate PCE router

path-request

Syntax

path-request [lsp-type {rsvp-p2p}] [dest ip-address] [detail]

Context

show>router>pcep>pcc

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

This command displays PCEP PCC path request information.

Parameters

lsp-type

Specifies the type of LSP to display. The only available option is RSVP-TE point-to-point LSPs.

ip-address

Specifies the destination IPv4 address to display.

Values a.b.c.d

detail

Keyword to display detailed path request information.

Output

The following output is an example of PCEP PCC path request information, and [Table 44: Output fields: Path request](#) describes the output fields.

Sample output

```
*A:ZBG# show router pcep pcc path-request
=====
PCEP Path Computation Client (PCC) Path Computation Request (PCReq) Info
=====
Request ID      : 4                Message State    : sent-for-compute
Tunnel ID       : 2                Extended Tunnel Id : 10.20.1.3
```



```
LSP ID      : 62468      LSP Type      : rsvp-p2p
LSP Name    : test_lsp::fully_loose
Source Address : 10.20.1.3   Destination Address: 10.20.1.1
SVEC Id     : 4           LSP Bandwidth   : 0
=====
*A:ZBG#
```

Table 44: Output fields: Path request

Label	Description
Request ID	Displays the PCEP PCC path request identifier
Tunnel ID	Displays the tunnel identifier for the LSP
LSP ID	Displays the LSP identifier
LSP Name	Displays the configured LSP name
Source Address	Displays the source IP address of the LSP
SVEC Id	Displays the synchronization vector identifier
Message State	Displays the current state of the request
Extended Tunnel Id	Displays the expanded tunnel identifier for the LSP
LSP Type	Displays the type of the LSP
Destination Address	Displays the destination IP address of the LSP
LSP Bandwidth	Displays the bandwidth of the LSP

peer

Syntax

peer [ip-address] [detail]

Context

show>router>pcep>pcc

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

This command displays PCEP PCC peer information.

Parameters

ip-address

Specifies a peer IPv4 address to display.

Values a.b.c.d

detail

Keyword to display detailed peer information.

Output

The following output is an example of a PCEP PCC peer information, and [Table 45: Output fields: PCC peer](#) describes the output fields.

Sample output

```

*A:ZBG>show>router>pcep>pcc# peer detail
=====
PCEP Path Computation Client (PCC) Peer Info
=====
IP Address           : 10.10.10.11
Admin Status         : Down           Oper Status           : Down
Peer Capabilities    : (Not Specified)
Speaker ID           : (Undefined)
Sync State           : not-initialized Peer Overloaded       : False
Session Establish Time: 0d 00:00:00
Oper Keepalive       : N/A           Oper DeadTimer        : N/A
Session Setup Count   : 0           Session Setup Fail Count: 0
-----
Statistics Information
-----

```

	Sent	Received
PC Request Message	0	0
PC Reply Message	0	0
PC Error Message	0	0
PC Notification Message	0	0
PC Keepalive Message	0	0
PC Update Message	0	0
PC Report Message	0	0
Path Report	0	0
Path Request	0	0

```

=====
*A:ZBG>show>router>pcep>pcc#

```

Table 45: Output fields: PCC peer

Label	Description
IP Address	Displays the IP address of the PCC peer
Admin Status	Displays the administrative status of the PCC peer
Oper Status	Displays the operational status of the PCC peer
Peer Capabilities	Displays the PCEP capabilities of the PCC peer

Label	Description
Speaker ID	Displays the IP address of the PCC peer speaker
Sync State	Displays the synchronization state of the PCC peer speaker
Peer Overloaded	Displays whether the PCC peer is overloaded
Session Establish Time	Displays the length of time since the PCEP session was established
Oper Keepalive	Displays the operational value for the PCC peer keepalive timer
Oper DeadTimer	Displays the operational value for the PCC peer dead timer
Session Setup Count	Displays the number of times that the PCEP session has been set up
Session Setup Fail Count	Displays the number of times that the PCEP session failed to be set up
Statistics Information	
PC Request Message	Displays the number of path computation (PC) request messages sent the PCC peer and received from the PCC peer
PC Reply Message	Displays the number of PC reply messages sent to the PCC peer and received from the PCC peer
PC Error Message	Displays the number of PC error messages sent to the PCC peer and received from the PCC peer
PC Notification Message	Displays the number of PC notification messages sent to the PCC peer and received from the PCC peer
PC Keepalive Message	Displays the number of PC keepalive messages sent to the PCC peer and received from the PCC peer
PC Update Message	Displays the number of PC update messages sent to the PCC peer and received from the PCC peer
PC Report Message	Displays the number of PC report messages sent to the PCC peer and received from the PCC peer
Path Report	Displays the number of path reports sent to the PCC peer and received from the PCC peer
Path Request	Displays the path requests sent to the PCC peer and received from the PCC peer

status

Syntax
status

Context
show>router>pcep>pcc

Platforms
7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description
This command displays PCEP PCC status information.

Output
The following output is an example of a PCEP PCC status information, and [Table 46: Output fields: PCC status](#) describes the output fields.

Sample output

```
*A:ZBG>show>router>pcep>pcc# status
=====
Path Computation Element Protocol (PCEP) Path Computation Client (PCC) Info
=====
Admin Status           : Down           Oper Status           : Down
Unknown Msg Limit      : 10 msg/min
Keepalive Interval     : 50 seconds    DeadTimer Interval    : 150 seconds
Capabilities List       : stateful-delegate stateful-pce rsvp-path
Address                 : 10.10.10.10
Report Path Constraints: True
-----
PCEP Path Computation Client (PCC) Peer Info
-----
Peer                   Admin State/Oper State Oper Keepalive/Oper DeadTimer
-----
10.10.10.11            Down/Down                Not-Applicable/Not-Applicable
-----
*A:ZBG>show>router>pcep>pcc#
```

Table 46: Output fields: PCC status

Label	Description
Admin Status	Displays the administrative status of the PCC
Oper Status	Displays the operational status of the PCC
Unknown Msg Limit	Displays the maximum rate of unknown messages that can be received on a PCEP session

Label	Description
Keepalive Interval	Displays the specified keepalive interval for the PCEP session
DeadTimer Interval	Displays the specified dead time interval for the PCEP session
Capabilities List	Displays the capabilities list for the PCEP session
Address	Displays the local IP address of the PCEP speaker
Report Path Constraints	Displays whether to include LSP path constraints in the PCE report messages sent from the PCC to a PCE
PCEP Path Computation Client (PCC) Peer Info	
Peer	Displays the IP address of the PCC peer
Admin State/Oper State	Displays the administrative and operational states of the PCC peer
Oper Keepalive/Oper Dead Timer	Displays the operational keepalive and dead timer intervals of the PCC peer

4.7.2.3 Tools commands

pcc

Syntax

pcc lsp [*plsp-id plsp-id*]

pcc lsp lsp-type *lsp-type* [*tunnel-id tunnel-id*]

Context

tools>dump>router>pcep

Platforms

7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12

Description

This command displays PCEP PCC LSP information.

Parameters

lsp

Keyword to display LSP information.

plsp-id

Specifies the ID of a PCC LSP. Only information for the PCC LSP with the specified ID is displayed.

Values 1 to 1048575

lsp-type

Specifies an LSP type. Only information for LSPs matching the specified type is displayed.

Values rsvp-p2p

tunnel-id

Specifies a tunnel ID. Only information for the tunnel with the specified ID is displayed.

Values 1 to 65535

5 Standards and protocol support



Note:

The information provided in this chapter is subject to change without notice and may not apply to all platforms.

Nokia assumes no responsibility for inaccuracies.

The following conventions are used in this section:

- T(A,N) indicates 7210 SAS-T in both Access-uplink mode and Network mode. Similarly, T(N) 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.

5.1 BGP

draft-ietf-idr-add-paths-04, Advertisement of Multiple Paths in BGP is supported on T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-ietf-sidr-origin-validation-signaling-04, BGP Prefix Origin Validation State Extended Community is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1772, Application of the Border Gateway Protocol in the Internet is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1997, BGP Communities Attribute is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2385, Protection of BGP Sessions via the TCP MD5 Signature Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2439, BGP Route Flap Damping is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2545, Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2858, Multiprotocol Extensions for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2918, Route Refresh Capability for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3107, Carrying Label Information in BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3392, Capabilities Advertisement with BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4271, A Border Gateway Protocol 4 (BGP-4) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4360, BGP Extended Communities Attribute is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4364, BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4456, BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Only R6 and R12 supports RR server functionality. Rest of the platforms support only client function.

RFC 4659, BGP/MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4684, Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4760, Multiprotocol Extensions for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4798, Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4893, BGP Support for Four-octet AS Number Space is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5004, Avoid BGP Best Path Transitions from One External to Another is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5291, Outbound Route Filtering Capability for BGP-4 is supported on K12, K30, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5668, 4-Octet AS Specific BGP Extended Community is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6811, Prefix Origin Validation is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

5.2 Certificate Management

RFC 5280, Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile 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 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.

5.3 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

5.4 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-/10GE standalone mode only.

draft-ietf-bess-evpn-vpws-14, Virtual Private Wire Service support in Ethernet VPN is supported on Mxp

5.5 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.

5.6 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-rrp-unified-spec-02, Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6 is supported on Mxp

RFC 768, User Datagram Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 793, Transmission Control Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 854, TELNET Protocol Specifications is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 951, Bootstrap Protocol (BOOTP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1034, Domain Names - Concepts and Facilities is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1035, Domain Names - Implementation and Specification is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1350, The TFTP Protocol (revision 2) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1534, Interoperation between DHCP and BOOTP is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1542, Clarifications and Extensions for the Bootstrap Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2131, Dynamic Host Configuration Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2132, DHCP Options and BOOTP Vendor Extensions is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2347, TFTP Option Extension is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2348, TFTP Blocksize Option is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2349, TFTP Timeout Interval and Transfer Size Options is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2428, FTP Extensions for IPv6 and NATs is supported on D, Dxp, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2865, Remote Authentication Dial In User Service (RADIUS) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2866, RADIUS Accounting is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3046, DHCP Relay Agent Information Option (Option 82) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3596, DNS Extensions to Support IP version 6 is supported on D, Dxp, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3768, Virtual Router Redundancy Protocol (VRRP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4250, The Secure Shell (SSH) Protocol Assigned Numbers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4251, The Secure Shell (SSH) Protocol Architecture is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4252, The Secure Shell (SSH) Authentication Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

All 7210 platforms support password and public key 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

5.7 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 is not supported.

RFC 4607, Source-Specific Multicast for IP is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4608, Source-Specific Protocol Independent Multicast in 232/8 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4610, Anycast-RP Using Protocol Independent Multicast (PIM) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5059, Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5384, The Protocol Independent Multicast (PIM) Join Attribute Format is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6513, Multicast in MPLS/BGP IP VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6514, BGP Encodings and Procedures for Multicast in MPLS/IP VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6515, IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6625, Wildcards in Multicast VPN Auto-Discover Routes is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6826, Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 7246, Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 7385, IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

5.8 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.

5.9 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

5.10 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.

5.11 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

5.12 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

5.13 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

5.14 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

5.15 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.

5.16 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

5.17 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

5.18 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

5.19 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

5.20 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

5.21 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

5.22 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.

5.23 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

5.24 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](#)