



7705 Service Aggregation Router

Release 26.10.R1

MPLS Guide

3HE 24956 AAAB TQZZA

Edition: 01

October 2026

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

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

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

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

No part of this document may be modified.

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

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

© 2026 Nokia.

Table of contents

List of tables.....	15
List of figures.....	18
1 Preface.....	21
1.1 Audience.....	21
1.2 Technical support.....	21
2 7705 SAR MPLS configuration process.....	22
3 MPLS and RSVP-TE.....	23
3.1 Overview.....	23
3.2 MPLS.....	24
3.2.1 Traffic engineering for MPLS.....	24
3.2.1.1 TE metric and IGP metric.....	25
3.2.2 MPLS label stack.....	25
3.2.2.1 Label values.....	26
3.2.2.2 Reserved label blocks.....	28
3.2.3 MPLS entropy labels.....	28
3.2.3.1 Overview of entropy labels.....	28
3.2.3.2 Inserting and processing the entropy label.....	31
3.2.3.3 Entropy label on OAM packets.....	32
3.2.3.4 Segment routing entropy label and IPsec, ESPI hashing, and NGE.....	32
3.2.3.5 Entropy label configuration.....	32
3.2.4 Label edge and label switching routers.....	33
3.2.5 LSP types.....	34
3.3 RSVP and RSVP-TE.....	34
3.3.1 RSVP-TE overview.....	35
3.3.1.1 Using RSVP-TE for MPLS.....	36
3.3.1.2 RSVP-TE extensions for MPLS.....	36
3.3.1.3 Hello protocol.....	37
3.3.1.4 Authentication.....	37
3.3.1.5 Non-router ID addresses as destinations and hops.....	39
3.3.2 RSVP LSP and LDP FEC statistics.....	39

3.3.2.1	Configuring RSVP LSP statistics at ingress LER.....	40
3.3.2.2	Configuring RSVP LSP statistics at egress LER.....	41
3.3.2.3	Configuring LDP FEC statistics.....	42
3.4	RSVP-TE signaling.....	42
3.4.1	General attributes of RSVP-TE.....	43
3.4.1.1	Bidirectional forwarding detection.....	43
3.4.1.2	Timers.....	43
3.4.1.3	LSP resignal limit.....	44
3.4.1.4	RSVP-TE message pacing.....	44
3.4.1.5	RSVP-TE overhead refresh reduction.....	44
3.4.1.6	RSVP-TE reservation styles.....	45
3.4.1.7	Implicit null label.....	46
3.4.1.8	RSVP-TE entropy labels.....	46
3.5	LSP redundancy.....	46
3.5.1	Make-before-break procedures for LSP and path parameter configuration changes.....	47
3.5.2	Automatic creation of RSVP-TE LSPs.....	48
3.5.3	Automatic creation of RSVP-TE LSP mesh (auto-LSP).....	48
3.5.3.1	Multi-area and multi-instance support.....	49
3.5.3.2	Mesh LSP name encoding.....	50
3.5.4	Automatic creation of an RSVP-TE single-hop LSP.....	50
3.5.5	Automatic ABR selection for inter-area LSPs.....	51
3.5.5.1	Rerouting of inter-area LSPs.....	53
3.5.5.2	Behavior of MPLS options in inter-area LSPs.....	53
3.5.5.3	Inter-area LSP support of OSPF virtual links.....	54
3.5.6	ABR FRR protection for inter-area LSP.....	54
3.6	Preference option for standby secondary LSP paths.....	56
3.7	RSVP-TE FRR.....	56
3.7.1	FRR terminology.....	57
3.7.2	Bypass resignal timer.....	59
3.7.3	FRR behavior.....	59
3.7.4	Dynamic and manual bypass LSPs.....	60
3.7.4.1	Bypass LSP selection rules for the PLR.....	60
3.7.4.2	FRR node protection (facility backup).....	62
3.7.5	Admin group support on facility bypass backup LSPs.....	63
3.7.6	FRR over unnumbered interfaces.....	63
3.8	Shared risk link groups.....	64

3.8.1	SRLGs for secondary LSP paths.....	64
3.8.2	SRLGs for FRR LSP paths.....	64
3.8.3	Disjoint and non-disjoint paths.....	65
3.8.4	Enabling disjoint backup paths.....	65
3.9	RSVP-TE graceful shutdown.....	67
3.10	RSVP-TE support for unnumbered interfaces.....	67
3.11	PCEP support for RSVP-TE LSPs.....	69
3.12	Segment routing with traffic engineering (SR-TE).....	69
3.12.1	SR-TE support.....	70
3.12.2	SR-TE LSP instantiation.....	71
3.12.2.1	PCC-initiated and PCC-controlled LSPs.....	72
3.12.2.2	PCC-initiated and PCE-computed/controlled LSPs.....	75
3.12.3	SR-TE LSP path computation.....	77
3.12.3.1	Service and shortcut application SR-TE label stack check.....	77
3.12.4	SR-TE LSP path computation using local CSPF.....	79
3.12.4.1	Extending MPLS and TE database CSPF support to SR-TE LSPs.....	80
3.12.4.2	SR-TE specific TE-DB changes.....	81
3.12.4.3	SR-TE LSP specific CSPF changes.....	81
3.12.5	SR-TE LSP paths using explicit SIDs.....	86
3.12.6	SR-TE LSP protection.....	87
3.12.6.1	Local protection.....	89
3.12.6.2	End-to-end protection.....	89
3.12.7	Seamless BFD for SR-TE LSPs.....	89
3.12.7.1	Configuration of S-BFD on SR-TE LSPs.....	90
3.12.7.2	Support for BFD failure action with SR-TE LSPs.....	91
3.12.7.3	S-BFD operational considerations.....	92
3.12.8	Static route resolution using SR-TE LSPs.....	92
3.12.9	BGP label route resolution using SR-TE LSPs.....	93
3.12.10	Service packet forwarding using SR-TE LSPs.....	93
3.12.11	Data path support.....	94
3.12.11.1	SR-TE LSP metric and MTU settings.....	96
3.12.12	SR-TE entropy labels.....	97
3.13	Weighted ECMP by RSVP-TE or SR-TE LSPs.....	97
3.13.1	Weighted ECMP for VPRN services with autobind.....	98
3.13.2	Weighted ECMP for EVPN.....	99
3.13.3	Weighted ECMP for SDPs.....	99

3.14	MPLS service usage.....	99
3.14.1	Service destination points.....	99
3.15	MPLS and RSVP-TE configuration process overview.....	100
3.16	Configuration notes.....	100
3.17	Configuring MPLS and RSVP-TE with CLI.....	101
3.18	MPLS configuration overview.....	101
3.18.1	Router interface.....	101
3.18.1.1	E-LSP for differentiated services.....	101
3.18.2	Paths.....	102
3.18.3	LSPs.....	102
3.18.4	Pseudowires.....	102
3.18.5	Signaling protocol.....	102
3.19	Basic MPLS configuration.....	103
3.20	Common configuration tasks.....	104
3.20.1	Configuring global MPLS parameters.....	104
3.20.2	Configuring an MPLS interface.....	105
3.20.3	Configuring MPLS paths.....	106
3.20.4	Configuring an MPLS LSP.....	107
3.20.5	Configuring a static LSP.....	108
3.20.5.1	Configuring a fast-retry timer for static LSPs.....	108
3.20.6	Configuring manual bypass tunnels.....	109
3.20.7	Configuring RSVP-TE parameters and interfaces.....	110
3.20.8	Configuring RSVP-TE message pacing parameters.....	111
3.21	MPLS configuration management tasks.....	111
3.21.1	Deleting MPLS.....	112
3.21.2	Modifying MPLS parameters.....	112
3.21.3	Modifying an MPLS LSP.....	112
3.21.4	Modifying MPLS path parameters.....	112
3.21.5	Modifying MPLS static LSP parameters.....	113
3.21.6	Deleting an MPLS interface.....	114
3.22	RSVP-TE configuration management tasks.....	114
3.22.1	Modifying RSVP-TE parameters.....	114
3.22.2	Modifying RSVP-TE message pacing parameters.....	115
3.22.3	Deleting an interface from RSVP-TE.....	115
3.23	Configuring and operating SR-TE.....	116
3.23.1	SR-TE configuration prerequisites.....	116

3.23.2	SR-TE LSP configuration overview.....	117
3.23.3	Configuring path computation and control for SR-TE LSPs.....	117
3.23.3.1	Configuring path profile and group for PCC-initiated and PCE-computed/ controlled LSPs.....	118
3.23.4	Configuring SR-TE LSP label stack size.....	118
3.23.5	Configuring adjacency SID parameters.....	118
3.23.6	Configuring PCC-controlled, PCE-computed, and PCE-controlled SR-TE LSPs.....	119
3.24	MPLS and RSVP-TE command reference.....	122
3.24.1	Command hierarchies.....	122
3.24.1.1	MPLS commands.....	122
3.24.1.2	RSVP-TE commands.....	125
3.24.1.3	Show commands.....	126
3.24.1.4	Clear commands.....	127
3.24.1.5	Monitor commands.....	127
3.24.1.6	Debug commands.....	127
3.24.2	Command descriptions.....	129
3.24.2.1	Configuration commands (MPLS).....	129
3.24.2.2	Configuration commands (RSVP-TE).....	194
3.24.2.3	Show commands (MPLS).....	210
3.24.2.4	Show commands (MPLS-labels).....	250
3.24.2.5	Show commands (seamless BFD).....	254
3.24.2.6	Show commands (RSVP).....	257
3.24.2.7	Clear commands.....	268
3.24.2.8	Monitor commands.....	271
3.24.2.9	Debug commands.....	278
4	PCEP.....	291
4.1	Introduction to the Path Computation Element Communication Protocol.....	291
4.2	Path computation fallback for PCC-initiated LSPs.....	294
4.3	Base implementation of path computation elements (PCE).....	295
4.4	PCEP session establishment and maintenance.....	296
4.5	PCEP parameters.....	297
4.5.1	PCC configuration.....	298
4.5.2	Stateful PCE.....	298
4.5.3	PCEP extensions in support of SR-TE LSPs.....	300
4.5.4	LSP initiation.....	301

4.5.5	PCC-initiated and PCE-computed or PCE-controlled LSPs.....	302
4.6	PCEP over TLS.....	303
4.6.1	TLS handshake.....	304
4.6.2	PCEP session over TLS.....	304
4.7	PCEP support for RSVP-TE LSPs.....	305
4.7.1	RSVP-TE LSP configuration for a PCC router.....	305
4.7.2	Behavior of the LSP path update.....	306
4.7.2.1	Path update with empty ERO.....	306
4.7.3	Behavior of LSP MBB.....	307
4.7.3.1	PCC-controlled LSPs.....	307
4.7.3.2	PCE-computed LSPs.....	308
4.7.3.3	PCE-controlled LSPs.....	308
4.7.4	Behavior of secondary LSP paths.....	310
4.7.5	PCE path profile support.....	310
4.8	LSP path diversity and bidirectionality constraints with path profiles.....	311
4.9	PCEP association groups.....	312
4.9.1	Diversity association groups.....	314
4.9.2	Policy association groups.....	314
4.10	NSP and VSR-NRC PCE redundancy.....	314
4.10.1	Feature configuration.....	314
4.10.2	NSP ecosystem redundancy.....	315
4.10.2.1	NSP cluster redundancy.....	315
4.10.2.2	VSR-NRC redundancy.....	316
4.10.2.3	PCC behavior.....	317
4.11	Configuring RSVP-TE LSPs with PCEP using the CLI.....	318
4.11.1	PCEP on the PCE node and the PCC node.....	318
4.11.2	MPLS on the PCC node.....	320
4.12	PCEP configuration command reference.....	324
4.12.1	Command hierarchies.....	324
4.12.1.1	PCEP commands.....	324
4.12.1.2	Show commands.....	325
4.12.2	Command descriptions.....	326
4.12.2.1	PCEP commands.....	326
4.12.2.2	Show commands.....	338
5	ANYsec.....	347

5.1	ANYsec overview.....	347
5.2	ANYsec packet format.....	348
5.3	ANYsec encryption.....	349
5.3.1	Encryption algorithms.....	349
5.3.2	MPLS protocol support.....	349
5.3.3	Per-flow encryption.....	349
5.4	ANYsec and MACsec interaction.....	349
5.5	ANYsec and LAG and ECMP interaction.....	349
5.6	Inter-AS and inter-area solutions.....	350
5.7	ANYsec implementation design.....	350
5.8	ANYsec configuration guidelines.....	351
5.8.1	Configuring ANYsec connectivity association and PSK.....	351
5.8.2	Identifying and configuring ANYsec LSP.....	352
5.8.3	Configuring ANYsec MKA.....	354
5.8.4	Configuring ANYsec encryption SIDs.....	355
5.9	ANYsec OAM MPLS support.....	358
5.10	ANYsec service encryption.....	358
5.10.1	Encryption Groups.....	359
5.10.1.1	Shared EG between services.....	360
5.10.1.2	Unique EG between services.....	361
5.10.2	Security termination policy.....	361
5.10.3	Service encryption considerations.....	361
5.10.4	Inter-AS and intra-AS solutions.....	362
5.10.5	Key management.....	362
5.10.6	Packet format for services.....	364
5.11	ANYsec command reference.....	367
5.11.1	Command hierarchies.....	367
5.11.1.1	Configuration commands.....	367
5.11.1.2	Show commands.....	368
5.11.1.3	Clear commands.....	368
5.11.1.4	Debug commands.....	368
5.11.2	Command descriptions.....	369
5.11.2.1	Configuration commands.....	369
5.11.2.2	Show commands.....	381
5.11.2.3	Clear commands.....	386
5.11.2.4	Debug commands.....	386

6	Label Distribution Protocol.....	389
6.1	Label Distribution Protocol.....	389
6.1.1	LDP and MPLS.....	390
6.1.1.1	BFD for T-LDP.....	391
6.1.2	LDP architecture.....	391
6.1.3	LDP subsystem interrelationships.....	391
6.1.3.1	Memory manager and LDP.....	391
6.1.3.2	Label manager.....	392
6.1.3.3	LDP configuration.....	392
6.1.3.4	Logger.....	393
6.1.3.5	Service manager.....	393
6.1.4	Execution flow.....	393
6.1.4.1	Initialization.....	393
6.1.4.2	Session lifetime.....	393
6.1.5	Label exchange.....	394
6.1.5.1	Implicit null label.....	394
6.1.5.2	Other reasons for label actions.....	395
6.1.5.3	Cleanup.....	395
6.1.6	LDP filters.....	395
6.1.6.1	Per-LDP peer export policies.....	396
6.1.7	LDP FEC statistics.....	396
6.1.8	Multi-area and multi-instance extensions to LDP.....	396
6.1.9	ECMP support for LDP.....	397
6.1.9.1	Label operations.....	399
6.1.10	Graceful restart helper.....	399
6.1.11	Graceful handling of resource exhaustion.....	400
6.1.12	LDP support for unnumbered interfaces.....	400
6.1.13	LDP FRR.....	402
6.1.13.1	ECMP vs FRR.....	402
6.1.13.2	IGP shortcuts (RSVP-TE tunnels).....	403
6.1.13.3	LDP FRR configuration.....	403
6.1.14	LDP-to-segment routing stitching for IPv4 /32 prefixes (IS-IS).....	403
6.1.14.1	Stitching in the LDP-to-SR direction.....	405
6.1.14.2	Stitching in the SR-to-LDP direction.....	406
6.1.14.3	TTL propagation and ICMP tunneling.....	407

6.1.15	LDP FRR remote LFA and TI-LFA backup using an SR tunnel for IPv4 /32 prefixes (IS-IS).....	407
6.1.15.1	Feature behavior.....	408
6.1.16	TCP MD5 and enhanced TCP authentication.....	408
6.2	LDP point-to-multipoint support.....	409
6.2.1	LDP point-to-multipoint configuration.....	409
6.2.2	LDP point-to-multipoint protocol.....	409
6.2.3	Make-before-break (MBB).....	409
6.2.4	ECMP support.....	409
6.2.5	Inter-AS non-segmented mLDP.....	410
6.2.5.1	LDP recursive FEC process.....	410
6.2.5.2	OAM.....	413
6.2.5.3	ECMP support.....	414
6.2.6	Intra-AS non-segmented mLDP.....	415
6.2.6.1	ABR MoFRR for intra-AS.....	415
6.2.6.2	Interaction with an inter-AS non-segmented mLDP solution.....	415
6.3	mLDP fast upstream switchover.....	416
6.3.1	mLDP fast upstream switchover configuration.....	416
6.3.2	mLDP fast upstream switchover behavior.....	417
6.4	LDP IPv6.....	418
6.4.1	Link LDP.....	419
6.4.2	Targeted LDP.....	420
6.4.3	FEC resolution.....	420
6.4.4	LDP session capabilities.....	421
6.4.5	LDP adjacency capabilities.....	422
6.4.6	IP address and FEC distribution.....	423
6.4.7	IGP and static route synchronization with LDP.....	425
6.4.8	BFD operation.....	426
6.4.9	Services using SDP with an LDP IPv6 FEC.....	426
6.4.10	Mirror services.....	427
6.4.10.1	Configuration at mirror source node.....	427
6.4.10.2	Configuration at mirror destination node.....	427
6.4.11	OAM support with LDP IPv6.....	427
6.4.12	Interoperability.....	428
6.4.12.1	Interoperability with implementations compliant with draft-ietf-mpls-ldp-ipv6... 428	
6.4.12.2	Interoperability with implementations compliant with RFC 5036 for IPv4 LDP control plane only.....	429

6.4.13	Upgrading from IPv4 to IPv6.....	429
6.5	LDP process overview.....	430
6.6	Configuration notes.....	430
6.7	Configuring LDP with CLI.....	431
6.8	LDP configuration overview.....	431
6.9	Basic LDP configuration.....	431
6.10	Common configuration tasks.....	431
6.10.1	Enabling LDP.....	432
6.10.2	Configuring graceful restart helper parameters.....	432
6.10.3	Applying import and export policies.....	433
6.10.4	Configuring interface parameters.....	434
6.10.5	Specifying targeted session parameters.....	435
6.10.6	Specifying peer parameters.....	436
6.10.7	Configuring LDP support for multicast VPN.....	437
6.10.8	Configuring LDP support for LDP-to-SR stitching.....	438
6.10.9	Enabling LDP signaling and services.....	438
6.11	LDP configuration management tasks.....	439
6.11.1	Disabling LDP.....	439
6.11.2	Modifying targeted session parameters.....	439
6.11.3	Modifying interface parameters.....	440
6.12	LDP command reference.....	442
6.12.1	Command hierarchies.....	442
6.12.1.1	LDP commands.....	442
6.12.1.2	Show commands.....	444
6.12.1.3	Monitor commands.....	446
6.12.1.4	Clear commands.....	446
6.12.1.5	Debug commands.....	447
6.12.2	Command descriptions.....	448
6.12.2.1	Configuration commands.....	448
6.12.2.2	Show commands.....	479
6.12.2.3	Clear commands.....	536
6.12.2.4	Monitor commands.....	539
6.12.2.5	Debug commands.....	546
7	List of acronyms.....	552

8	Supported standards and protocols.....	579
8.1	Security standards.....	579
8.2	Telecom standards.....	579
8.3	Protocol support.....	580
8.3.1	ATM.....	580
8.3.2	BFD.....	580
8.3.3	BGP.....	581
8.3.4	DHCP/DHCPv6.....	581
8.3.5	Differentiated services.....	582
8.3.6	Digital data network management.....	582
8.3.7	ECMP.....	582
8.3.8	Ethernet VPN (EVPN).....	582
8.3.9	Frame relay.....	582
8.3.10	GRE.....	582
8.3.11	Internet protocol (IP) – version 4.....	583
8.3.12	Internet protocol (IP) – version 6.....	583
8.3.13	IPSec.....	583
8.3.14	IS-IS.....	584
8.3.15	LDP.....	585
8.3.16	LDP and IP FRR.....	585
8.3.17	MPLS.....	585
8.3.18	MPLS – OAM.....	586
8.3.19	Multicast.....	586
8.3.20	Network management.....	586
8.3.21	OSPF.....	588
8.3.22	OSPFv3.....	588
8.3.23	PPP.....	588
8.3.24	Pseudowires.....	589
8.3.25	RIP.....	589
8.3.26	RADIUS.....	589
8.3.27	RSVP-TE and FRR.....	589
8.3.28	Segment routing (SR).....	590
8.3.29	SONET/SDH.....	590
8.3.30	SSH.....	590
8.3.31	Synchronization.....	590

8.3.32	TACACS+.....	591
8.3.33	TLS.....	591
8.3.34	TWAMP.....	592
8.3.35	VPLS.....	592
8.3.36	VRRP.....	592
8.4	Proprietary MIBs.....	592

List of tables

Table 1: Configuration process.....	22
Table 2: Packet/label field description.....	26
Table 3: Ingress label values (pop labels).....	27
Table 4: Egress label values (push labels).....	27
Table 5: Summary of entropy label support.....	30
Table 6: FRR terminology.....	57
Table 7: Parameter values for frr-overhead.....	78
Table 8: Disabled and enabled options for bypass-only.....	109
Table 9: Router MPLS admin-group field descriptions.....	211
Table 10: Router MPLS bypass-tunnel field descriptions.....	212
Table 11: Router MPLS interface field descriptions.....	215
Table 12: Router MPLS LSP field descriptions.....	218
Table 13: MPLS LSP detail field descriptions.....	219
Table 14: Router MPLS LSP path detail field descriptions.....	223
Table 15: Router MPLS LSP path MBB field descriptions.....	226
Table 16: Router MPLS auto-LSP field descriptions.....	227
Table 17: Router MPLS LSP template field descriptions.....	232
Table 18: Router MPLS path field descriptions.....	235
Table 19: Router MPLS static LSP field descriptions.....	247
Table 20: Router MPLS status field descriptions.....	249
Table 21: Router MPLS-labels label field descriptions.....	251

Table 22: Router MPLS-labels label range field descriptions.....	252
Table 23: Router MPLS-labels summary field descriptions.....	254
Table 24: MPLS seamless BFD session field descriptions.....	255
Table 25: Router RSVP-TE interface field descriptions.....	258
Table 26: Router RSVP-TE interface detail field descriptions.....	259
Table 27: Router RSVP-TE interface statistics field descriptions.....	261
Table 28: Router RSVP-TE neighbor field descriptions.....	263
Table 29: Router RSVP-TE session field descriptions.....	265
Table 30: Router RSVP-TE statistics field descriptions.....	266
Table 31: Router RSVP-TE status field descriptions.....	267
Table 32: Base PCEP TLVs, objects, and messages.....	295
Table 33: PCEP stateful PCE extension TLVs, objects, and messages.....	298
Table 34: PCEP segment routing extension objects and TLVs.....	301
Table 35: PCEP path profile extension objects and TLVs.....	312
Table 36: PCEP PCC field descriptions.....	338
Table 37: PCEP diversity association field descriptions.....	342
Table 38: PCEP policy association field descriptions.....	343
Table 39: PCEP PCC peer field descriptions.....	344
Table 40: PCEP PCC status field descriptions.....	346
Table 41: MKA over IP field descriptions.....	382
Table 42: Tunnel encryption field descriptions.....	384
Table 43: Opaque types.....	412
Table 44: OAM functionality for option C.....	414

Table 45: Hello timeout factor default values.....	460
Table 46: Keepalive timeout factor default values.....	461
Table 47: LDP discovery field descriptions.....	481
Table 48: FEC-originate field descriptions.....	485
Table 49: LDP interface field descriptions.....	487
Table 50: LDP parameters field descriptions.....	489
Table 51: LDP session field descriptions.....	493
Table 52: LDP status field descriptions.....	496
Table 53: LDP targeted peer field descriptions.....	499
Table 54: LDP bindings field descriptions.....	505
Table 55: Acronyms.....	552

List of figures

Figure 1: Label structure.....	26
Figure 2: Label packet placement.....	26
Figure 3: Entropy label and load balancing.....	32
Figure 4: Establishing LSPs.....	35
Figure 5: LSP using RSVP-TE path setup.....	36
Figure 6: Automatic ABR selection for inter-area LSP.....	52
Figure 7: ABR protection using dynamic bypass LSP.....	55
Figure 8: Bypass tunnel node example.....	61
Figure 9: FRR node-protection example.....	63
Figure 10: Disjoint primary and secondary LSPs.....	66
Figure 11: Disjoint FRR bypass LSPs.....	67
Figure 12: Label stack reduction in a 3-tier ring topology.....	85
Figure 13: Label stack reduction in the presence of ECMP paths.....	86
Figure 14: SR-TE LSP label stack programming.....	96
Figure 15: MPLS and RSVP-TE configuration and implementation flow.....	100
Figure 16: Manual bypass tunnels.....	109
Figure 17: NSP functional modules.....	292
Figure 18: NRC-P architecture.....	293
Figure 19: PCEP session initialization.....	297
Figure 20: NSP ecosystem redundancy.....	315
Figure 21: Multi-level IS-IS topology in the NSP GUI.....	318

Figure 22: Primary and secondary RSVP-TE LSP paths in the NSP GUI.....	320
Figure 23: ANYsec MPLS encryption and decryption.....	347
Figure 24: ANYsec high-level implementation.....	348
Figure 25: ANYsec MPLS packet authentication.....	348
Figure 26: ANYsec packet authentication with EL and ELI.....	350
Figure 27: ANYsec encryption implementation design and requirements.....	351
Figure 28: ANYsec tunnel encryption implementation.....	354
Figure 29: ANYsec using MKA UDP port configuration.....	355
Figure 30: ANYsec double-encryption scenario.....	356
Figure 31: ANYsec encryption SID.....	357
Figure 32: Tunnel and service encryption.....	358
Figure 33: Two encrypted services both using the same tunnel and EG.....	360
Figure 34: Two encrypted services using the same tunnel and different EGs.....	361
Figure 35: Services Option A.....	362
Figure 36: MKA for encrypted service.....	363
Figure 37: SAK for encrypted service.....	364
Figure 38: VPRN1 RT 1:1 service.....	365
Figure 39: VPRN2 RT 2:2 service.....	365
Figure 40: Example 2: Encrypted service and SR tunnel with different EGs.....	366
Figure 41: VPRN1 RT 1:1 service.....	367
Figure 42: VPRN2 RT 1:1 service.....	367
Figure 43: LDP subsystem interrelationships.....	392
Figure 44: Stitching in the LDP-to-SR direction.....	405

Figure 45: MLDP FEC for single AS with transit IPv4 opaque.....	410
Figure 46: MLDP FEC for inter-AS with recursive opaque value.....	411
Figure 47: Non-VPN mLDP with recursive opaque for inter-AS.....	411
Figure 48: ASBRs using EBGp without IGP.....	413
Figure 49: ECHO request target FEC stack TLV.....	413
Figure 50: ECMP support.....	414
Figure 51: Intra-AS non-segmented topology.....	415
Figure 52: MLDP fast upstream switchover.....	417
Figure 53: LDP adjacency and session over an IPv6 interface.....	419
Figure 54: LDP IPv6 address and FEC distribution procedure.....	424
Figure 55: LDP IPv6 address and FEC distribution procedure.....	425
Figure 56: Smooth management transition from IPv4 to IPv6.....	430
Figure 57: LDP configuration and implementation.....	430

1 Preface

This guide describes the services and protocol support provided by the 7705 SAR and presents examples to configure and implement MPLS (RSVP-TE and LDP) protocols and PCEP protocols.

This guide is organized into functional chapters and provides concepts and descriptions of the implementation flow, as well as Command Line Interface (CLI) syntax and command usage.



Note: This manual generically covers Release 26.x content and may contain some content that will be released in later maintenance loads. See the 7705 SAR 26.x.Rx Software Release Notes, part number 3HE24971000xTQZZA, for information about features supported in each load of the Release 26.x software.



Note: As of Release 23.4, software support for the following hardware has been deprecated:

- 8-port Ethernet Adapter card, version 2 (a8-ethv2) (3HE02776)
- 12-port Serial Data Interface card, version 1 (a12-sdi) (3HE03391)
- 7705 SAR-W (3HE07349)

These components are no longer recognized in the release.

If information about any of the above components is required, please see the applicable installation guides in Release 22.10.

1.1 Audience

This guide is intended for network administrators who are responsible for configuring the 7705 SAR routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Concepts described in this guide include the following:

- multiprotocol label switching (MPLS)
- Resource Reservation Protocol for Traffic Engineering (RSVP-TE)
- Label Distribution Protocol (LDP)

1.2 Technical support

If you purchased a service agreement for your 7705 SAR router and related products from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance. If you purchased a Nokia service agreement, follow this link to contact a Nokia support representative and to access product manuals and documentation updates:

[Product Support Portal](#)

2 7705 SAR MPLS configuration process

The following table lists the tasks that are required to configure MPLS, RSVP-TE, and LDP protocols.

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

Table 1: Configuration process

Area	Task/description	Chapter
Protocol configuration	Configure MPLS parameters	MPLS
	Configure RSVP-TE parameters	RSVP and RSVP-TE
	Configure PCEP parameters	PCEP
	Configure LDP parameters	Label Distribution Protocol
Reference	List of security and telecom standards, supported protocols, and proprietary MIBs	Supported standards and protocols

3 MPLS and RSVP-TE

This chapter provides information required to configure multiprotocol label switching (MPLS) and Resource Reservation Protocol for Traffic Engineering (RSVP-TE) for the 7705 SAR. For information about dynamic LSPs with LDP, see the chapter [Label Distribution Protocol](#).

Topics in this chapter include:

- [Overview](#)
- [MPLS](#)
- [RSVP and RSVP-TE](#)
- [RSVP-TE signaling](#)
- [LSP redundancy](#)
- [Preference option for standby secondary LSP paths](#)
- [RSVP-TE FRR](#)
- [Shared risk link groups](#)
- [RSVP-TE graceful shutdown](#)
- [RSVP-TE support for unnumbered interfaces](#)
- [PCEP support for RSVP-TE LSPs](#)
- [Segment routing with traffic engineering \(SR-TE\)](#)
- [Weighted ECMP by RSVP-TE or SR-TE LSPs](#)
- [MPLS service usage](#)
- [MPLS and RSVP-TE configuration process overview](#)
- [Configuration notes](#)
- [Configuring MPLS and RSVP-TE with CLI](#)
- [MPLS and RSVP-TE command reference](#)

3.1 Overview

The 7705 SAR provides MPLS technology using static LSPs, RSVP-TE for traffic-engineered signaled routing of LSPs and LDP for non-traffic-engineered signaled routing of LSPs. A network operator may choose to use any combination of static LSPs, RSVP-TE, and LDP to establish paths for services. RSVP-TE and LDP are considered to be Layer 2.5 protocols.

The 7705 SAR can be used as an ingress and egress label edge router (iLER and eLER), and as a transit router. A transit router is also referred to as a label switching router (LSR).

OSPF and IS-IS are the interior gateway protocols with traffic engineering extensions (IGP-TE) available to the 7705 SAR. These are the Layer 3 protocols. Typically, one or the other of these gateway protocols will be in use in the network. Whichever protocol is the chosen gateway protocol, it must be working in

order for LDP or RSVP-TE to function. These Layer 3 protocols identify the next hop, which is information needed by the Layer 2.5 protocols (LDP or RSVP-TE) in order to assign labels.

In addition, the 7705 SAR provides link and node redundancy protection through LSP redundancy and fast reroute (FRR) features.

The LSP redundancy and FRR features have the ability to take shared risk link groups (SRLGs) into consideration when the constrained shortest path first (CSPF) algorithm is used to determine an alternate LSP. The selection of a route is determined by the IGP-TE protocol. The added constraints imposed by SRLGs and CSPF ensure that the redundant route selected will be unique from the principal route (route being protected); that is, it will use physical equipment that is different from the equipment that carries the principal route. CSPF will constrain the alternate route to be the shortest possible alternative route. There may be more than one alternative route.

3.2 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 independent of any routing protocol but is considered multiprotocol because it works with protocols such as IP, ATM, Ethernet, and circuit emulation.

This section contains the following topics:

- [Traffic engineering for MPLS](#)
- [MPLS label stack](#)
- [MPLS entropy labels](#)
- [Label edge and label switching routers](#)
- [LSP types](#)

3.2.1 Traffic engineering for MPLS

Without traffic engineering (TE), routers route traffic according to the shortest path first (SPF) algorithm, disregarding congestion or packet types.

With traffic engineering, network traffic is routed efficiently to maximize throughput and minimize delay. Traffic engineering facilitates traffic flows to be mapped to the destination through a less-congested path 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 (bidirectional) 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 (from next hop to next hop) 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 using SPF to reach the destination IP address.

3.2.1.1 TE metric and IGP metric

When the TE metric is selected for an LSP, the shortest path computation will select an LSP path based on the TE metric constraints instead of the IGP metric (for OSPF and IS-IS), which is the default metric. The user configures the TE metric under the **router>mpls>interface** context and the IGP metric under the **router>ospf>area> interface** context (for OSPF) and the **router>isis>if>level** context (for IS-IS). 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 the IGP protocols. For more information about the OSPF and IS-IS routing protocols, see the 7705 SAR Routing Protocols Guide.

Typically, the TE metric is used to allow constrained shortest path first (CSPF) to represent a dual TE topology for the purpose of computing LSP paths, where one TE topology is based on the RSVP-TE database and the other is based on the IGP-TE database.

An LSP dedicated to 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 amount of delay, or combined delay and jitter, 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 to 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 value as required.

When the use of the TE metric is enabled for an LSP, the CSPF process will first eliminate all links in the network topology that do not meet the constraints specified for the LSP path; the constraints include bandwidth, admin-groups, and hop limit. CSPF will then run the SPF algorithm 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. The TE metric is only used in CSPF computations for MPLS paths and not in the regular SPF computation for IP reachability.

Operational metrics of LSPs that use the TE metric in CSPF path calculations can be overridden with the user-configured administrative LSP metric.

3.2.2 MPLS label stack

Routers that support MPLS are known as label edge routers (LERs) and label switching routers (LSRs). MPLS requires a set of procedures to enhance network layer packets with label stacks, which turns them into labeled packets. To initiate, transmit, or terminate a labeled packet on a particular data link, an LER or 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 (that is, remove the top label), or swap the label and push one or more labels onto 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 other labels may have been above it in the past or that 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 structure of a label and the table describes the fields.

Figure 1: Label structure





Table 2: 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 time-to-live 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



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.

3.2.2.1 Label values

The 7705 SAR uses RSVP-TE and LDP protocols for label forwarding, For packet-based services such as VLL, the 7705 SAR uses T-LDP for signaling PW labels between peer nodes.

Packets traveling along an LSP are identified by the packet label, which is the 20-bit, unsigned integer (see [Label edge and label switching routers](#)). The range is 0 through 1 048 575. Label values 0 to 15 are reserved and are defined below:

- A value of 0 represents the IPv4 Explicit NULL label. This label value is legal only at the bottom of the label stack if the label stack is immediately followed by an IPv4 header, in which case the packet forwarding is based on the IPv4 header. If the IPv4 Explicit NULL label is not at the bottom of the label stack, then the packet forwarding is based on the subsequent label.
- 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. Since this label cannot be 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. An LER advertises this when it is requesting penultimate hop popping (PHP) and expecting unlabeled packets. The label value 3 should never appear in the label stack.
- A value of 7 represents the entropy label indicator (ELI). The ELI is a special-purpose MPLS label that indicates that the entropy label (EL) follows it in the stack.
- Values 4 through 6 and 8 through 15 are reserved for future use.

The following table lists the label ranges available for use by ingress labels (pop labels).

Table 3: Ingress label values (pop labels)

Label values	Description
16 through 31	Reserved for future use
32 through 1023	Available for static outer LSP tunnel label assignment
1024 through 2047	Reserved for future use
2048 through 18 431 ¹	Statically assigned for services (inner pseudowire label)
32 768 through 131 071	Dynamically assigned for both MPLS and services
131 072 through 1 048 575	Reserved for future use

Note:

1. In addition, users can define part of the dynamic label range from 18 432 to 131 071 to be the range of labels for the segment routing global block (SRBG).

The following table lists the label ranges available for use by egress labels (push labels).

Table 4: Egress label values (push labels)

Label values	Description
16 through 1 048 575	Can be used for static LSP tunnel and static PW labels
16 through 1 048 575	Can be dynamically assigned for both MPLS tunnel labels and PW labels

3.2.2.2 Reserved label blocks

Reserved label blocks are sets of labels that are reserved for allocation by various applications. These blocks are separate from existing label ranges such as the static labels range. Labels are reserved from the dynamic label range and are not tied to the bottom of the label range. Applications can use the labels in a reserved block as local segment identifiers (SIDs). The reserved block is called a segment routing local block (SRLB). Labels are reserved from the dynamic label range. Up to four reserved label blocks can be configured on a system. Only one range can be configured for each block.

Use the following CLI to configure a reserved label block:

```
config>router>mpls-labels>reserved-label-block name>start-label start-value end-label end-value
```

Reserved label blocks apply to both OSPF and IS-IS segment routing (SR). To ensure that the label allocation process is serialized, label allocation within a reserved label range is managed by the SR module instead of the label manager. This means that any application needing to assign a label from a reserved range must request it from the SR module.

3.2.3 MPLS entropy labels

This section contains information about the following topics:

- [Overview of entropy labels](#)
- [Inserting and processing the entropy label](#)
- [Entropy label on OAM packets](#)
- [Segment routing entropy label and IPSec, ESPI hashing, and NGE](#)
- [Entropy label configuration](#)

3.2.3.1 Overview of entropy labels

The 7705 SAR supports MPLS entropy labels on RSVP-TE and SR-TE LSPs, as per RFC 6790. The entropy label provides greater granularity for load balancing on an LSR where load balancing is typically based on the MPLS label stack.

The ability of a node to receive and process an entropy label for an LSP is signaled using capability signaling (referred to as entropy label capability (ELC)). Entropy labels are supported on RSVP-TE and SR-TE tunnels.

Inserting an entropy label adds two labels in the MPLS label stack: the entropy label itself and the entropy label indicator (ELI).

The entropy label is inserted directly below the tunnel label and closest to the service payload that has advertised entropy label capability (which may be above the bottom of the stack). The value of the entropy label is calculated at the iLER and is based on a hash of the packet payload header content and other system parameters at ingress. For more information about hashing inputs, see the "Per-Flow Hashing" section in the 7705 SAR Interface Configuration Guide.

The ELI is inserted by the iLER. The ELI is a special-purpose MPLS label (value = 7) that indicates that the entropy label is the next label in the stack.

Entropy label capability is advertised at the tunnel level by the far-end node (eLER). This capability can be advertised for an RSVP-TE FEC or an SR-TE tunnel on IS-IS or OSPF. Capability signaling is not

supported for point-to-multipoint LSPs, BGP tunnels, or LDP FECs. An LSR used for RSVP-TE and SR-TE tunnels will pass the entropy label capability signal from the downstream LSP segment to upstream peers. However, earlier releases that do not support entropy label functionality will pass the capability flag transparently, without altering the value.

The insertion of an entropy label by the upstream LER on a tunnel enabled for entropy label capability is enabled on a per-service basis. The entropy label is only inserted if the downstream peer has signaled entropy label support. The upstream LER only inserts a single entropy label, even if multiple LSP labels exist in a label stack.

The 7705 SAR supports the entropy label feature for the following services:

- Cpipe, Epipe, and lpipe access to spoke SDP
- Cpipe, Epipe, and lpipe spoke SDP to spoke SDP (vc-switching)
- VPLS SAP to VPLS spoke SDP or mesh SDP
- VPLS spoke SDP to VPLS spoke SDP
- VPRN for RSVP-TE
- R-VPLS
- IGP shortcut
- IS-IS for SR-TE
- OSPF for SR-TE

Entropy label capability on RSVP-TE LSPs is enabled on the eLER using the **config>router>rsvp>entropy-label-capability** command.

At the iLER, the insertion of the entropy label into the label stack is enabled using the **entropy-label** command under the service, mesh SDP, or spoke SDP context or under the **config>router>isis (or ospf)>segment-routing** context for SR-TE LSPs.

The entropy label requires the insertion of two additional labels in the label stack. In some cases, this may result in an unsupported label stack depth or large changes in the label stack depth during the lifetime of an LSP (for example, due to switching from a primary path with entropy label capability enabled to a secondary path for which the far end has not signaled entropy label capability).

The **entropy-label** command under the **config>router>mpls** and **config>router>mpls>lsp** contexts provides local control at the head end of an LSP over whether the entropy label is inserted on an LSP by overriding the entropy label capability signaled from the far-end LER, and control over how the additional label stack depth is accounted for. This allows the user to avoid entropy label insertion where there is a risk of the label stack depth becoming too great.

For entropy labels that are supported on LDP tunnels with remote-LFA protection (that is, for **rsvp-shortcut**), only loop-free alternate protect (**lfa-protect**) and LFA (**lfa-only**) are allowed.

Support of entropy labels over RSVP-TE and SR-TE tunnels are the only valid options, except when the 7705 SAR is the LER node with BGP labeled unicast (BGP-LU) tunnels. A 7705 SAR in an LER role can push and pop an entropy label for Epipe and VPLS services with a BGP-LU tunnel riding over an RSVP-TE LSP. Conversely, a 7705 SAR does not support being in an ABR or ASBR role with BGP-LU. The following table lists entropy label support on the 7705 SAR.

Table 5: Summary of entropy label support

Service	RSVP-TE	SR-TE
Epipe	Yes	Yes
lpipe	Yes	Yes
Cpipe	Yes	Yes
Apipe, Fpipe, Hpipe	No	No
VPRN (MP-BGP)	Yes	Yes
VPRN (Layer 3 spoke SDP)	Yes	Yes
IES (Layer 3 spoke SDP)	Yes	Yes
VPLS SDP (spoke/mesh SDP)	Yes	Yes
LDP over IGP shortcut (RSVP)	Yes	N/A
IGP shortcut (SR)	N/A	No
LDP FRR over RSVP	Yes	N/A
LDP stitching over SR (SR to LDP)	N/A	Yes ¹
LDP stitching over SR (LDP to SR)	No	No
BGP LU ²	Yes	Yes
SR	No	Yes
EVPN VPLS	Yes	Yes
EVPN Epipe	Yes	Yes
R-VPLS	Yes	Yes
IGP shortcut	Yes	No
SR FRR over TI-LFA or R-LFA	N/A	Yes
Static route with tunnel next hop	Yes	Yes

Notes:

1. On the SR segment because the SR head end injects the entropy label.
2. For services that support entropy label.

3.2.3.2 Inserting and processing the entropy label

This section contains inserting and processing information for the following node types:

- [Ingress LER](#)
- [LSR](#)
- [Egress LER](#)

3.2.3.2.1 Ingress LER

The procedures at the iLER are as specified in section 4.2 of RFC 6790. In general, the router inserts an entropy label into the label stack if the downstream node for the LSP tunnel has signaled support for entropy label and the entropy label is enabled for the particular service.

RFC 6790 specifies that the iLER can insert several entropy labels in the label stack where the LSP hierarchy exists, one for each LSP in the hierarchy. However, this could result in unreasonably large label stacks. Therefore, when there are multiple LSPs in a hierarchy (for example, LDP over RSVP-TE), the router only inserts a single EL/ELI pair within the innermost LSP label closest to the service payload that has advertised entropy label capability.

The router inserts an entropy label on a tunnel that is entropy label-capable when the service has entropy label enabled, even if an implicit or explicit NULL label has been signaled by the downstream LSR or LER. This ensures consistent behavior and ensures that the entropy label value as determined by the iLER is maintained where a tunnel with an implicit NULL label is stitched at a downstream LSR.

3.2.3.2.2 LSR

If an LSR is configured for load balancing and an entropy label is found in the label stack, the LSR will take the entropy label into account in the hashing algorithm as follows:

- **label-only** – the entropy label is used as input to the hash routine and the rest of the label stack is ignored.
- **label-ip** – the entropy label and the IP packet are used as input to the hash routine and the rest of the label stack is ignored.

If penultimate hop popping (PHP) has been requested by a next-hop LER, the LSR will retain any entropy label found immediately below the tunnel label that is to be popped. The system will retain and use the entropy label information as input to the local hash routine if an applicable LSR load-balancing mode has been configured.

For more information about LSR load balancing, see the *LSR Hashing* section in the 7705 SAR Interface Configuration Guide.

3.2.3.2.3 Egress LER

At an eLER, if an ELI and entropy label are detected in the label stack, both the ELI and entropy label are popped and the packet processed as normal. This occurs whether or not the system has signaled entropy label capability.

If an ELI is popped that has the bottom of stack (BoS) bit set, the system will discard the packet.

3.2.3.3 Entropy label on OAM packets

Service OAM packets also include an entropy label and ELI if entropy label capability is signaled for the corresponding tunnel and entropy label is enabled for the service. The EL/ELI pair is inserted at the same level in the label stack as it is in user data packets; that is, within the innermost LSP label context closest to the service payload that has advertised entropy label capability. The EL/ELI pair will therefore always reside at a different level in the label stack from special-purpose labels related to the service payload (for example, the router alert label).

OAM packets at the LSP level, such as LSP ping and LSP trace, do not have the EL/ELI pair inserted.

3.2.3.4 Segment routing entropy label and IPsec, ESPI hashing, and NGE

Segment routing with entropy label can be used with IPsec and NGE services and with ESPI hashing, as listed below:

- IPsec and segment routing entropy label
 - IPsec over BGP 3107 over segment routing with entropy label
 - IPsec over static route over segment routing with entropy label
 - VLL over GRE over IPsec over BGP 3107 over segment routing with entropy label
 - VLL over GRE over IPsec over static route over segment routing with entropy label
- ESPI hashing GRT/VRN
- NGE
 - VLL, VPLS, and VRN NGE interaction with entropy label


3.2.3.5 Entropy label configuration

The following figure illustrates the use of entropy labels at the service level.

The iLER has entropy label enabled under an applicable service context and the eLER has entropy label capability enabled. The iLER inserts the ELI and the EL into the label stack. The entropy label value is based on the service ID for point-to-point Layer 2 services.

At the LSR, if hashing is enabled, the LSR recognizes the ELI and uses the entropy label value as the hash result. If the **entropy-label** command had been disabled at the iLER, the LSR would not find the ELI and would default to hashing based on the label stack, if applicable.

Figure 3: Entropy label and load balancing



26222

At the ingress LER:

```
config>service>cpipe>spoke-sdp>entropy-label
config>service>epipe>spoke-sdp>entropy-label
config>service>ipipe>spoke-sdp>entropy-label
config>service>vpls>spoke-sdp>entropy-label
config>service>vpls>mesh-sdp>entropy-label
config>service>vprn>entropy-label
config>service>vprn>interface>spoke-sdp>entropy-label
config>router>isis>segment-routing>entropy-label
config>router>ospf>segment-routing>entropy-label
```

At the egress LER:

```
config>router>entropy-label
config>router>rsvp>entropy-label-capability
config>router>mpls>lsp>entropy-label
config>router>isis>entropy-label>override-tunnel-elc
config>router>ospf>entropy-label>override-tunnel-elc
```

The **per-service-hashing** command and the **l4-load-balancing**, **teid-load-balancing**, and **spi-load-balancing** commands are mutually exclusive.

For IP traffic, use the **l4-load-balancing** command. For IP traffic with mobile payload, use the **teid-load-balancing** and/or the **spi-load-balancing** command.

3.2.4 Label edge and label switching routers

A 7705 SAR performs different functions based on its position in an LSP—ingress, egress, or transit—as described in the following list:

- ingress label edge router (iLER) – the router at the beginning of an LSP is the iLER. The ingress router encapsulates packets with an MPLS header and forwards the packets to the next router along the path. An LSP can only have one ingress router.
- label switching router (LSR) – an LSR can be any intermediate router in the LSP between the ingress and egress routers, swapping the incoming label with the outgoing MPLS label and forwarding the MPLS packets it receives to the next router in the LSP. An LSP can have 0 to 253 transit routers.
- egress label edge router (eLER) – the router at the end of an LSP is the 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. An LSP can have only one egress router. The ingress and egress routers in an LSP cannot be the same router.

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

Constrained-path LSPs are signaled and are confined to one Interior Gateway Protocol (IGP) area. These LSPs cannot cross an autonomous system (AS) boundary.

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

3.2.5 LSP types

The following LSP types are supported:

- static LSPs – a static LSP specifies a static path. All routers that the LSP traverses must be configured manually with labels. No RSVP-TE or LDP signaling is required. Static LSPs are discussed in this chapter.
- signaled LSPs – LSPs are set up using the RSVP-TE or LDP signaling protocol. 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. RSVP-TE is discussed in this chapter, and LDP is discussed in [Label Distribution Protocol](#).

There are two types of signaled LSP:

- 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 other routers (loose). This enables the user to control how the path is set up. Explicit-path LSPs are similar to static LSPs but require less configuration. See [RSVP and RSVP-TE](#). An explicit path that has not specified any hops will follow the IGP route.
- constrained-path LSPs – for 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 that satisfies the constraints for the LSP. In turn, CSPF relies on the topology database provided by an extended IGP such as OSPF or IS-IS.

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

If fast reroute (FRR) is configured, the ingress router signals the downstream routers so that each downstream router can preconfigure a detour route for the LSP that will be used if there is a failure on the original LSP. If a downstream router does not support FRR, the request is ignored and the router continues to support the original LSP. This can cause some of the detour routes to fail, but the original LSP is not impacted. For more information about FRR, see [RSVP-TE FRR](#).

No bandwidth is reserved for the reroute 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 establishing the backup LSP. The following warning message is displayed:

"The fast reroute bandwidth command is not supported in this release."

3.3 RSVP and RSVP-TE

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 operational state to provide the requested service. In general, RSVP requests result in resources reserved in each node along the data path.

The Resource Reservation Protocol for Traffic Engineering (RSVP-TE) is an extended version of RSVP for MPLS. RSVP-TE uses traffic engineering extensions to support automatic signaling of LSPs. MPLS uses RSVP-TE to set up traffic-engineered LSPs. See [RSVP-TE extensions for MPLS](#) for more information.

3.3.1 RSVP-TE overview


RSVP-TE requests resources for simplex (unidirectional) flows. Therefore, RSVP-TE 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-TE is a signaling protocol, not a routing protocol. RSVP-TE operates with unicast and multicast routing protocols. Routing protocols determine where packets are forwarded. RSVP-TE consults local routing tables to relay RSVP-TE messages.

RSVP-TE uses two message types to set up LSPs, PATH and RESV. [Figure 4: Establishing LSPs](#) depicts 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 forwarding equivalence class (FEC) for which label bindings are desired. PATH messages are used to signal and request the 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 4: Establishing LSPs



20120

The following figure displays an example of an LSP path set up using RSVP-TE. The ingress label edge router (iLER 1) transmits an RSVP-TE 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 5: LSP using RSVP-TE path setup



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

- explicit route object (ERO) – forces the RSVP-TE PATH message 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
- session attribute object – controls the path setup priority, holding priority, and local rerouting features

Upon receiving a PATH message containing a label request object, the eLER transmits an 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.

3.3.1.1 Using RSVP-TE for MPLS

Hosts and routers that support both MPLS and RSVP-TE can associate labels with RSVP-TE flows. When MPLS and RSVP-TE are combined, the definition of a flow can be made more flexible. Once 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 Forwarding Equivalence Class (FEC) that defines the RSVP-TE flow.

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

3.3.1.2 RSVP-TE extensions for MPLS

The RSVP-TE extensions enable MPLS to support the creation of explicitly routed LSPs, with or without resource reservation. Several of the features enabled by these extensions were implemented to meet the requirements for traffic engineering over MPLS, which enables the creation of traffic trunks with specific characteristics. None of the TE extensions result in backward compatibility problems with traditional RSVP implementations.

To run properly, the traffic engineering capabilities of RSVP-TE require an underlying TE-enabled IGP routing protocol. The 7705 SAR supports OSPF and IS-IS with TE extensions.

Routing protocols make it possible to advertise the constraints imposed over various links in the network. For example, in order for the nodes in a network to choose the best link for signaling a tunnel, the capacity of a particular link and the amount of reservable capacity must be advertised by the IGP. RSVP-TE makes use of these constraints to request the setup of a path or LSP that traverses only those links that are part of an administrative group (admin groups are described in the following list). Therefore, both RSVP-TE and the IGP-TE (that is, OSPF-TE or IS-IS-TE for the 7705 SAR) must be enabled and running simultaneously.

The following TE capabilities are supported:

- hop limit – the hop limit is the maximum number of LSRs that a given LSP can traverse, including the ingress and the egress LERs. Typically, the hop limit is used to control the maximum delay time for mission-critical traffic such as voice traffic.

The hop limit applies to the primary LSP, any backup LSPs, and LSPs configured to be used in FRR situations.

- admin groups – administrative groups provide a way to define which LSR nodes should be included or excluded while signaling an LSP. For example, it may be desirable to avoid some nodes or links that are known to be used heavily from being included in the path of an LSP, or to include a specific LSR node to ensure that a newly signaled RSVP-TE tunnel traverses that LSR node.

Administrative groups apply to both primary and secondary LSPs. They are defined under the **config>router>if-attribute** context, and are applied at the MPLS interface level, as well as at the LSP and the primary and secondary LSP levels through **include** and **exclude** commands.

- bandwidth – the bandwidth capability (supported by RSVP-TE), is similar to the Connection Admission Control (CAC) function in ATM. During the establishment phase of RSVP-TE, the LSP PATH message contains the bandwidth reservation request. If the requested capacity is available, the RESV message confirms the reservation request. The amount of reserved bandwidth stated in the request is deducted from the amount of reservable bandwidth for each link over which the LSP traverses.

The bandwidth capability applies to both primary and secondary LSPs, and LSPs configured to be used in FRR situations.

3.3.1.3 Hello protocol

The Hello protocol detects the loss of a neighbor node (node failure detection) or the reset of a neighbor's RSVP-TE 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 LERs. 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-TE 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.

3.3.1.4 Authentication

Protocol authentication protects against malicious attacks on the communications between routing protocol neighbors. These attacks could either disrupt communications or inject incorrect routing information into the systems routing table. The use of authentication keys can help to protect routing protocols from these types of attacks.

All RSVP-TE protocol exchanges can be authenticated. This guarantees that only trusted routers can participate in autonomous system routing.

Authentication must be explicitly configured and can be done using two separate mechanisms:

- configuration of an explicit authentication key and algorithm using the **authentication-key** command
- configuration of an authentication keychain using the **auth-keychain** command

Either the **authentication-key** command or the **auth-keychain** command can be used by RSVP-TE, but both cannot be supported at the same time. If both commands are configured, the **auth-keychain** configuration will be applied and the **authentication-key** command will be ignored.

By default, authentication is not enabled on an interface.

3.3.1.4.1 Authentication key

When enabled on an RSVP-TE interface with the **authentication-key** command, 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
- the key used with the authentication algorithm
- the lifetime of the key. A key is a user-generated key using third-party software or hardware. The value is entered as a static string into the CLI configuration of the RSVP interface. The key will continue to be valid until it is removed from that RSVP interface.
- the source address of the sending system
- the 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 that also contains a Flags field, a Key Identifier field, and a Sequence Number field. The RSVP sender complies with 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.

If a point of local repair (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 merge point (MP) node, then the message is discarded since there is no security association with the next-next-hop MP node.

The 7705 SAR MD5 implementation does not support the authentication challenge procedures in RFC 2747.

3.3.1.4.2 Authentication keychains

The keychain mechanism allows for the creation of keys used to authenticate RSVP-TE communications. Each keychain entry defines the authentication attributes to be used in authenticating RSVP-TE messages from remote peers or neighbors; the entry must include at least one key entry to be valid. The keychain mechanism also allows authentication keys to be changed without affecting the state of the RSVP-TE adjacencies and supports stronger authentication algorithms than plaintext and MD5.

Keychains are configured in the **config>system>security>keychain** context. For more information about configuring keychains, see the 7705 SAR System Management Guide, "TCP Enhanced Authentication and Keychain Authentication".

The keychain is then associated with an RSVP-TE interface with the **auth-keychain** command.

For a key entry to be valid, it must include a valid key, the current system clock value must be within the begin and end time of the key entry, and the algorithm specified must be supported by RSVP-TE.

RSVP-TE supports the following authentication algorithms:

- HMAC-MD5
- HMAC-SHA-1-96
- HMAC-SHA-1
- HMAC-SHA-256

Keychain errors are handled as follows:

- If a keychain exists but there are no active key entries with an authentication type that matches the type supported by RSVP-TE, inbound RSVP-TE packets will not be authenticated and will be discarded and no outbound RSVP-TE packets will be sent.
- If a keychain exists but the last key entry has expired, a log entry will be raised indicating that all keychain entries have expired.

RSVP-TE requires that the protocol continue to authenticate inbound and outbound traffic using the last valid authentication key.

3.3.1.5 Non-router ID addresses as destinations and hops

The address of a configured loopback interface, other than the router ID, can be used as the destination of an RSVP LSP. For a constrained-path LSP, CSPF searches for the best path that matches the constraints across all areas or levels of the IGP where this address is reachable. If the address is the router ID of the destination node, then CSPF selects the best path across all areas or levels of the IGP for that router ID, regardless of which area or level the router ID is reachable for as an interface.

The address of a loopback interface other than the router ID can also be configured as a hop in the LSP path hop definition. If the hop is "strict" and corresponds to the router ID of the node, the CSPF path may use any TE-enabled link to the downstream node based on best cost. If the hop is "strict" and does not correspond to the router ID of the node, CSPF will fail.

3.3.2 RSVP LSP and LDP FEC statistics

RSVP LSP and LDP FEC statistics allow operators to monitor traffic being forwarded between any two PE routers and for all services using an RSVP or LDP SDP. If the LSP is used as a shortcut to transport BGP LU, VPRN traffic over MP-BGP or IGP prefixes, statistics are collected for these IP packets being forwarded.

The following statistics are collected for each RSVP LSP or LDP FEC:

- per forwarding class forwarded in-profile packet count
- per forwarding class forwarded in-profile byte count
- per forwarding class forwarded out-of-profile packet count

- per forwarding class forwarded out-of-profile byte count

For an RSVP LSP, these counters are available for the egress data path at the ingress LER and for the ingress data path at the egress LER.

For an LDP FEC, these counters are available for the egress data path at the ingress LER and LSR. Because an ingress LER is also potentially an LSR for an LDP FEC, combined egress data path statistics are provided whenever applicable.

OAM packets that are forwarded using LSP encapsulation, such as LSP ping and LSP trace, are also included in the above counters.

Dropped packets and bytes for an RSVP LSP or LDP FEC are not counted on the ingress LER.

Octet counters are for the entire frame and include the label stack and Layer 2 header and padding, similar to existing MPLS interface counters. For that reason, ingress and egress octet counters for an RSVP LSP may differ slightly if the type of interface or encapsulation is different (POS, Ethernet null, or Ethernet dot1q).

RSVP LSP and LDP FEC statistics counters can be retrieved by:

- using the CLI **show** command for the RSVP LSP or the LDP FEC
- using the CLI **monitor** command applied to a specific RSVP LSP or LDP FEC
- using an SNMPv3 interface to query the MIB
- accessing an accounting file if statistics collection is enabled with the default or a user-specified accounting policy for the RSVP LSP or LDP FEC

RSVP LSP and LDP FEC statistics counters are not saved to an accounting file unless statistics collection is enabled and the specific RSVP LSP or LDP FEC statistics collection record is included in the default accounting policy or in a user-defined accounting policy using the following commands:

- RSVP LSP ingress data path counters

```
config>router>mpls>ingress-statistics>lsp>collect-stats
config>router>mpls>ingress-statistics>lsp>accounting-policy policy-id
```
- RSVP LSP egress data path counters

```
config>router>mpls>lsp>egress-statistics>collect-stats
config>router>mpls>lsp>egress-statistics>accounting-policy policy-id
```
- LDP FEC egress data path counters

```
config>router>ldp>egress-statistics>fec-prefix>collect-stats
config>router>ldp>egress-statistics>fec-prefix>accounting-policy policy-id
```

3.3.2.1 Configuring RSVP LSP statistics at ingress LER

At the ingress LER, statistics are configured in the egress data path of an originating LSP with the **config>router>mpls>lsp>egress-statistics** command in the LSP configuration at the head-end node. Statistics collection in the egress data path is enabled after the user executes the **no shutdown** command in the **egress-statistics** context. By default, this function is in a shutdown state.

Statistics cannot be configured if the LSP name contains a colon (:), which is used as a field separator by the ingress LER for encoding the LSP and path names into the RSVP Session Name field in the Session_Attribute object.

The **no** form of the **egress-statistics** command disables statistics collection in the egress data path and removes the accounting policy association from the RSVP LSP. Users can choose to disable statistics in the egress data path while keeping the accounting policy association of the RSVP LSP with the **config>router>mpls>lsp>egress-statistics shutdown** command.

The same set of counters are updated for packets forwarded over any path of the LSP. In the steady state, counters are updated for packets forwarded over the active path of the LSP. The active path can be the primary path, one of the secondary paths, the FRR detour path, or the FRR bypass path when the head-end node is also the PLR.

The LSP counters are maintained over the lifetime of the LSP as long as statistics are enabled. The user can clear the counters with the **clear>router>mpls>lsp-egress-stats [lsp-name]** command.

LSP statistics are not collected on a dynamic or static bypass tunnel. LSP egress statistics are also not collected if the head-end node is also the penultimate-popping hop (PHP) node for a single-hop LSP using an implicit null label. However, if any label is pushed onto the label stack, for example, the Layer 2 or Layer 3 service label, the egress statistics are counted for the LSP even if the transport MPLS label is not present.

When a hierarchy of LSPs is in use, statistics collection on the outermost label corresponding to the tunneling LSP and on the inner labels, corresponding to the tunneled LSPs, are mutually exclusive. The outermost label takes precedence. Consequently, when the user enables statistics collection on an RSVP LSP that is also used for tunneling LDP FECs with the LDP over RSVP shortcut, statistics will be collected on the RSVP LSP only. No statistics are collected for an LDP FEC tunneled over this RSVP LSP even if the user enabled statistics collection on the FEC. When the user disables statistics collection on the RSVP LSP, statistics collection, if enabled, will be performed on the tunneled LDP FEC.

LSP statistics are not collected on static LSPs. Auto-LSP templates do not support LSP statistics collection.

3.3.2.2 Configuring RSVP LSP statistics at egress LER

At the egress LER, statistics are configured in the ingress data path of a terminating LSP by entering the LSP name, along with the ingress LER system interface address, with the **config>router>mpls>ingress-statistics>lsp lsp-name sender ip-address** command. Statistics collection is enabled in the ingress data path after the user executes the **no shutdown** command in the **ingress-statistics** context. By default, this function is in a shutdown state.

The LSP name must correspond to the name configured by the user at the ingress LER. Statistics cannot be configured if the LSP name contains a colon (:), which is used as a field separator by the ingress LER for encoding the LSP and path names into the RSVP Session Name field in the Session_Attribute object.

The **no** form of the **ingress-statistics** command disables statistics collection in the ingress data path and removes the accounting policy association from the RSVP LSP. Users can choose to disable statistics in the ingress data path while keeping the accounting policy association of the RSVP LSP with the **config>router>mpls>ingress-statistics>lsp>shutdown** command.

The same set of counters are updated for packets received over any path of the LSP. In the steady state, the counters are updated for packets received over the active path of the LSP. The active path can be the primary path, one of the secondary paths, the FRR detour path, or the FRR bypass path when the tail-end node is also the MP.

The LSP counters are maintained over the lifetime of the LSP as long as statistics are enabled. The user can clear the counters with the **clear>router>mpls>lsp-ingress-stats ip-address lsp lsp-name** command.

When a hierarchy of LSPs is in use, statistics collection on the outermost label corresponding to the tunneling LSP and on the inner labels, corresponding to the tunneled LSPs, are mutually exclusive. The outermost label takes precedence.

Because ingress data path statistics are not supported for an LDP FEC, there are no statistics collected for an LDP FEC, however if the LDP FEC is tunneled over an RSVP shortcut LSP that has LSP ingress statistics configured, the statistics are collected for the RSVP LSP.

The user can enable statistics collection on a manual bypass LSP terminating on the egress LER. However, all LSPs whose primary path is protected by the manual bypass will not collect statistics when they activate forwarding over the manual bypass. If the user disables statistics collection on the manual bypass LSP, statistics collection, if enabled, is continued on the protected LSP when the bypass LSP is activated.

A flag in the output of the **show** command for the LSP statistics will indicate if there were no path state blocks (PSBs) that matched the specified LSP name at any specific point in time. This could be due to the absence of the RSVP session or to the presence of a session type that is not supported; for example, the LSP name matched a point-to-multipoint LSP. The counters will show all zero values, which could otherwise be confused with an LSP with a valid matched PSB that is not receiving packets.

3.3.2.3 Configuring LDP FEC statistics

At the ingress LER and LSR, statistics collection is configured in the egress data path of an LDP FEC by specifying the FEC prefix with the **config>router>ldp>egress-statistics>fec-prefix** command. Statistics collection is enabled in the egress data path after the user executes the **no shutdown** command under the **egress-statistics** context. By default, this function is in a shutdown state.

The **no** form of the **egress-statistics** command disables statistics collection in the egress data path and removes the accounting policy association from the LDP FEC. Users can choose to disable statistics in the egress data path while keeping the accounting policy association of the LDP FEC with the **config>router>ldp>egress-statistics>fec-prefix>shutdown** command.

Statistics collection applies to prefix FECs imported from both LDP neighbors and T-LDP neighbors.

The egress data path counters are updated for both originating and transit packets. Originating packets may be service packets or IP user and control packets forwarded either as BGP LU over LDP FEC or as VPRN traffic (MP-BGP) over LDP FEC or simply over LDP FEC IGP shortcut. Transit packets of the FEC are label-switched on this node.

When ECMP is enabled and multiple paths exist for a FEC, the same set of counters is updated for each packet forwarded over any of the ECMP links for as long as this FEC is active.

The LDP FEC counters are maintained over the lifetime of the FEC as long as statistics are enabled. The user can clear the counters with the **clear>router>ldp>fec-egress-statistics** command.

For more information about LDP FEC statistics commands, see [LDP commands](#).

3.4 RSVP-TE signaling

RSVP-TE-based signaling provides a means to establish tunnels dynamically.

RSVP-TE uses the downstream on demand (DOD) label distribution mode, sending PATH messages from the ingress LER node to the egress LER and RESV messages in the reverse direction. DOD label distribution is a router's response to an explicit request from another router for label binding information.

The DOD mode is in contrast to LDP on the 7705 SAR, which uses the Downstream Unsolicited (DU) label distribution mode for both PWs and LSPs. A router in DU mode will distribute label bindings to another router that has not explicitly requested the label bindings.

RSVP-TE signaling is supported when the 7705 SAR is deployed as an LER and as an LSR. When used as an LER, the 7705 SAR uses RSVP-TE signaling to set up constrained paths because only the LER knows all the constraints imposed on the LSP. When used as an LSR, the 7705 SAR uses RSVP-TE to interpret the RSVP-TE messages (including all the constraints).

With RSVP-TE, users can choose which services and PWs may use a particular LSP. One-to-one or many-to-one scenarios for binding PWs to RSVP-TE LSPs is supported, which is similar to binding PWs to static LSPs. Furthermore, each RSVP-TE LSP can be configured with its own set of attributes and constraints.

3.4.1 General attributes of RSVP-TE

The following general attributes of RSVP-TE on the 7705 SAR are supported:

- [Bidirectional forwarding detection](#)
- [Timers](#)
- [LSP resignal limit](#)
- [RSVP-TE message pacing](#)
- [RSVP-TE overhead refresh reduction](#)
- [RSVP-TE reservation styles](#)
- [Implicit null label](#)
- [RSVP-TE entropy labels](#)

3.4.1.1 Bidirectional forwarding detection

Bidirectional forwarding detection (BFD) is supported on the 7705 SAR. With BFD for RSVP-TE, an RSVP-TE enabled link is registered with the BFD state machine and if a failure occurs, the RSVP-TE interface is taken out of service. The BFD implementation on the 7705 SAR works on a hop-by-hop basis, and if BFD detects a link failure, only the two directly connected MPLS nodes are aware of that failure. If the node that detects the link failure is an LSR node, it generates PATH-ERR messages to the originators (the LER nodes) of the failing LSPs. If FRR is configured, the detecting node takes corrective action itself. See [LSP redundancy](#) and [RSVP-TE FRR](#) for more information about these topics.

3.4.1.2 Timers

The following timers are implemented to ensure the successful operation of RSVP-TE:

- `bypass-resignal-timer` – defines the time between the global reoptimization of all dynamic bypass RSVP-TE LSPs. For more information, see [Bypass resignal timer](#).
- `hold-timer` – defines the amount of time before an LSP is brought up and is in service, which provides protection against unreliable nodes and links
- `resignal-timer` – used in conjunction with the route optimization process, especially after a reroute has occurred. If the newly computed path for an LSP has a better metric than the currently recorded hop list, an attempt is made to resignal that LSP, and if the attempt is successful, a make-before-break

switchover occurs. If the attempt to resignal an LSP fails, the LSP continues to use the existing path and another resignal attempt is made the next time the timer expires.

When the resignal timer expires, a trap and syslog message are generated.

- **retry-timer** – defines a period of time before a resignal attempt is made after an LSP failure. This delay time protects network resources against excessive signaling overhead.

3.4.1.3 LSP resignal limit

When an LSP fails, an LER node tries to resignal it. The following limit can be configured:

- **retry-limit** – the retry limit defines the number of ressignaling attempts in order to conserve the resources of the nodes in the network. There could be a serious loss of capacity due to a link failure where an infinite number of retries generate unnecessary message overhead.

3.4.1.4 RSVP-TE message pacing

RSVP-TE message pacing provides a means to limit the overwhelming number of RSVP-TE signaling messages that can occur in large MPLS networks during node failures. RSVP-TE message pacing allows the messages to be sent in timed intervals.

To protect nodes from receiving too many messages, the following message pacing parameters can be configured:

- **msg-pacing** – message pacing can be enabled or disabled
- **max-burst** – maximum burst defines the number of RSVP-TE messages that can be sent in the specified period of time
- **period** – period defines the interval of time used in conjunction with the max-burst parameter to send message pacing RSVP-TE messages

Message pacing needs to be enabled on all the nodes in a network to ensure the efficient operation of tier-1 nodes. Message pacing affects the number of RSVP-TE messages that a particular node can generate, not the number of messages it can receive. Therefore, each node must be paced at a rate that allows the most loaded MPLS nodes to keep up with the number of messages they receive.



Note: Typically, a tier-1 node is an aggregator of tier-2 node transmissions, which is an aggregator of tier-3 node transmissions. Tier-1 nodes are often installed at an MTSO, while tier-3 nodes are often installed at cell sites.

3.4.1.5 RSVP-TE overhead refresh reduction

RFC 2961, *RSVP Refresh Overhead Reduction Extensions*, defines enhancements to the RSVP-TE signaling protocol that reduce refresh overhead, which are in addition to the message pacing function.

These extensions are:

- **RSVP-TE message bundling** – RSVP-TE message bundling reduces the total number of RSVP-TE messages by aggregating the status information of multiple LSPs into a single RSVP-TE PDU. The 7705 SAR supports the receipt and processing of bundled RSVP-TE messages but not the transmission of bundled messages as specified in RFC 2961, section 3.3.

- **reliable message delivery** – reliable message delivery extends RSVP-TE to support MESSAGE_ACK. Each RSVP-TE PDU has a unique message-id for sequence tracking purposes. When an RSVP-TE message arrives, the recipient acknowledges the reception of the specific message-id (this is similar to TCP ACK messages). Lost PDUs can be detected and re-sent with this method, which helps reduce the refresh rate because there are two endpoints tracking the received/lost messages.
- **summary refresh** – the summary refresh capability uses a single message-id list to replace many individual refresh messages and sends negative ACKs (NACKs) for any message-id that cannot be matched (verified). The summary refresh capability reduces the number of message exchanges and message processing between peers. It does not reduce the amount of soft state stored in the node. The term soft state refers to the control state in hosts and routers that will expire if not refreshed within a specified amount of time (see RFC 2205 for information about soft state).

These capabilities can be enabled on a per-RSVP-TE interface basis and are referred to collectively as “refresh overhead reduction extensions”. When **refresh-reduction** is enabled on a 7705 SAR RSVP-TE interface, the node indicates this to its peer by setting a refresh-reduction-capable bit in the flags field of the common RSVP-TE header. If both peers of an RSVP-TE interface set this bit, all three of the capabilities listed above can be used. The node monitors the setting of this bit in received RSVP-TE messages from the peer on the interface. If the bit is cleared, the node stops sending summary refresh messages. If a peer did not set the refresh-reduction-capable bit, a 7705 SAR node does not attempt to send summary refresh messages.

Also, reliable delivery of RSVP-TE messages over the RSVP-TE interface can be enabled using the **reliable-delivery** option.

3.4.1.6 RSVP-TE reservation styles

LSPs can be signaled with explicit reservation styles for the reservation of resources, such as bandwidth. A reservation style describes a set of attributes for a reservation, including the sharing attributes and sender selection attributes. The style information is part of the LSP configuration. The 7705 SAR supports two reservation styles:

- **fixed filter (FF)** – the fixed filter (FF) reservation style specifies an explicit list of senders and a distinct reservation for each of them. Each sender has a dedicated reservation that is not shared with other senders. Each sender is identified by an IP address and a local identification number, the LSP ID. Because each sender has its own reservation, a unique label and a separate LSP can be constructed for each sender-receiver pair. For traditional RSVP applications, the FF reservation style is ideal for a video distribution application in which each channel (or source) requires a separate pipe for each of the individual video streams.
- **shared explicit (SE)** – the shared explicit (SE) 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.

If the FRR option is enabled for the LSP and the facility FRR method is selected at the head-end node, only the SE reservation style is allowed. If a 7705 SAR 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 backup method supports both FF and SE styles.

3.4.1.7 Implicit null label

The implicit null label option enables an eLER to receive MPLS packets from the previous-hop LSR without the outer LSP label.

The implicit null label is included in RESV messages sent by the eLER to the previous-hop LSR. When the implicit null label is signaled to the LSR, it pops the outer label before sending the MPLS packet to the eLER; this is known as penultimate hop popping.

The implicit null label option can be enabled for all RSVP-TE interfaces and for all RSVP-TE LSPs for which the router is the eLER by using the **implicit-null-label** command in the **config>router>rsvp** context.

RSVP-TE must be shut down before this command can be used.

The implicit null label option can also be enabled or disabled on a specific RSVP-TE interface, overriding the RSVP-TE level configuration, by using the **implicit-null-label {enable | disable}** command in the **config>router>rsvp>interface** context.

The implicit null label is enabled for all LSPs for which the router is the eLER and for which the PATH message is received from the previous-hop LSR over the RSVP-TE interface.

With facility backup, if the eLER is also the merge point (MP) node, the incoming interface for the PATH refresh message over the bypass tunnel dictates whether the packet will use the implicit null label. Similarly, with one-to-one backup, if the eLER is also the detour merge point (DMP) node, the incoming interface for the PATH refresh message over the detour LSP dictates whether the packet will use the implicit null label.

The RSVP-TE interface must be shut down before this command can be used.

3.4.1.8 RSVP-TE entropy labels

The 7705 SAR supports entropy labels as described in [MPLS entropy labels](#).

3.5 LSP redundancy

Each primary LSP can be protected by up to two secondary LSPs. When the LER detects a primary LSP failure, it signals its secondary LSPs, if any have been configured, and automatically switches to the first one that is available. LSP redundancy supports shared risk link groups (SRLG). See [Shared risk link groups](#) for more information about SRLG.

LSP redundancy differs from the fast reroute (FRR) feature in that LSP redundancy is controlled by the LER that initiated the LSP, whereas FRR uses the node that detects the failure to take recovery action. This means that LSP redundancy takes longer to reroute traffic than FRR because failure messages need to traverse multiple hops to reach the LER and activate LSP redundancy, whereas an FRR-configured node responds immediately to bypass the failed node or link. See [RSVP-TE FRR](#) for more information about FRR.

The following parameters can be configured for primary and secondary LSPs:

- **bandwidth** – the amount of bandwidth needed for the secondary LSP can be reserved and can be any value; it does not need to be identical to the value reserved by the primary LSP. Bandwidth reservation can be set to 0, which is equivalent to reserving no bandwidth.

- inclusion and exclusion of nodes – by including or excluding certain nodes, you can ensure that the primary and secondary LSPs do not traverse the same nodes and therefore ensure successful recovery. Each secondary LSP can have its own list of included and excluded nodes.
- hop limit – the hop limit is the maximum number of LSRs that a secondary LSP can traverse, including the ingress and egress LERs.
- standby (secondary LSPs only) – when a secondary LSP is configured for standby mode, it is signaled immediately and is ready to take over traffic the moment the LER learns of a primary LSP failure. This mode is also called hot-standby mode.

When a secondary LSP is not in standby mode, then it is only signaled when the primary LSP fails. If there is more than one secondary LSP, they are all signaled at the same time (upon detection of a primary LSP failure) and the first one to come up is used.

If a **path-preference** priority value is configured for standby secondary LSP paths, the standby secondary LSP configured with the highest path priority becomes the active LSP when the primary LSP fails.

3.5.1 Make-before-break procedures for LSP and path parameter configuration changes

The make-before-break (MBB) procedure allows an LSP to switch from an existing working path to a new path without interrupting service. The MBB procedure does this by first signaling the new path when it is operationally up, having the ingress LER move the traffic to the new path, and then allowing the ingress LER to tear down the original path.

The MBB procedure is invoked during the following operations:

- timer-based and manual resignal of an LSP path
- fast reroute (FRR) global revertive procedures
- traffic engineering (TE) graceful shutdown procedures
- update of the secondary path due to an update to the primary path SRLG
- LSP primary or secondary path name change
- LSP or path configuration parameter change

MBB procedure coverage has been extended to most of the other LSP-level and path-level parameters as follows:

- including or excluding admin groups at the LSP and path levels
- enabling or disabling the LSP-level CSPF option: **lsp>path-computation-method local-cspf** or **lsp-template>cspf**
- enabling or disabling LSP-level metric type parameters when the CSPF option is enabled: **lsp>metric-type** or **lsp-template>cspf [use-te-metric]**
- enabling or disabling the LSP-level hop-limit option in the fast-reroute context
- enabling the LSP-level least-fill option
- enabling or disabling the LSP-level **adspec** option
- changing between node-protect and no node-protect (link-protect) values in the LSP-level fast-reroute option
- changing LSP primary or secondary path priority values (setup priority)
- changing the LSP-level and path-level hop-limit parameter values

- enabling or disabling primary or secondary path record or record-label options

The MBB procedure is not supported on a manual bypass LSP.

3.5.2 Automatic creation of RSVP-TE LSPs

Automatic creation of RSVP-TE LSPs enables the automated creation of point-to-point RSVP-TE LSPs within a single IGP IS-IS level or OSPF area that can subsequently be used by services and/or IGP shortcuts. The feature is divided into two modes: creation of an RSVP-TE LSP mesh, and creation of single-hop RSVP-TE LSPs.

When creating an RSVP-TE LSP mesh, the mesh can be full or partial, the extent of which is governed by a prefix list containing the system addresses of all nodes that should form part of the mesh. When using single-hop RSVP-TE LSPs, point-to-point LSPs are established to all directly connected neighbors.

The use of automatically created RSVP-TE LSPs avoids manual configuration of RSVP-TE LSP meshes. Even when provisioning tools are used to automatically provision these LSPs, automatic creation of a mesh still provides a benefit by avoiding increased configuration file sizes.

3.5.3 Automatic creation of RSVP-TE LSP mesh (auto-LSP)

This feature enables the automatic creation of an RSVP-TE point-to-point LSP to a destination node whose router ID matches a prefix in the specified peer prefix policy. This LSP type is referred to as an auto-created LSP mesh. To start the process of automatically creating an RSVP-TE LSP mesh, the user must create a route policy referencing a prefix list. This prefix list contains the system addresses of all nodes that are required to be in the mesh, and can be entered as a series of /32 addresses, or simply as a range.

After the route policy is created, the user must create an LSP template containing the common parameters that are used to establish all point-to-point LSPs within the mesh. The template must be created with the keyword **mesh-p2p**:

```
config>router>mpls>lsp-template template-name mesh-p2p
```

Upon creation of the template, CSPF is automatically enabled and cannot be disabled. The template must also reference a default path before it can be placed in a **no shutdown** state.

Next, the user must associate the LSP template with the previously defined route policy, and this is accomplished using the **auto-lsp lsp-template** command:

```
config>router>mpls>auto-lsp lsp-template template-name policy peer-prefix-policy
```

Once the **auto-lsp lsp-template** command is entered, the system starts the process of establishing the point-to-point LSPs. The prefixes defined in the prefix list are checked, and if a prefix corresponds to a router ID that is present in the Traffic Engineering (TE) database, the system instantiates a CSPF-computed primary path to that prefix using the parameters specified in the LSP template.

Multiple templates can be associated with the same or different peer prefix policies. Each application of an LSP template with a given prefix in the prefix list results in the instantiation of a single CSPF-computed LSP primary path using the LSP template parameters, as long as the prefix corresponds to a router ID for a node in the TE database. Auto-LSP does not support the automatic signaling of a secondary path for an LSP. If the signaling of multiple LSPs to the same destination node is required, a separate LSP template must be associated with a prefix list that contains the same destination node address. Each instantiated LSP will have a unique LSP ID and a unique tunnel ID.

The auto-created LSP is installed in the tunnel table manager (TTM) and is available to applications such as resolution of BGP label routes, and resolution of BGP, IGP, and static routes. The auto-created LSP can

also be used for autobinding by a VPRN service. The auto-created LSP cannot be used as a provisioned SDP for explicit binding by services.

The auto-created LSP mesh can be signaled over both numbered and unnumbered RSVP-TE interfaces.

Up to five peer prefix policies can be associated with an LSP template. Every time the user executes the **auto-lsp** command with the same or different prefix policy associations or changes the prefix policy associated with an LSP template, the system re-evaluates the prefix policy. The outcome of the re-evaluation indicates to MPLS whether an existing LSP must be torn down or a new LSP must be signaled to a destination address that is already in the TE database.

If a /32 prefix is added to or removed from a prefix list associated with an LSP template, or if a prefix range is expanded or narrowed, the prefix policy re-evaluation is performed. Whether the prefix list contains one or more specific /32 addresses or a range of addresses, MPLS requires an external trigger to instantiate an LSP to a node whose address matches an entry in the prefix list. The external trigger is when the router with a router ID matching an address in the prefix list appears in the TE database. The TE database provides the trigger to MPLS.

The user must perform a **no shutdown** of the template before it takes effect. When a template is in use, the user must shut down the template before changing any parameters except for those LSP parameters for which the change can be handled with the make-before-break (MBB) procedures (see [Make-before-break procedures for LSP and path parameter configuration changes](#)). When the template is shut down and parameters are added, removed, or modified, the existing instances of the LSP using this template are torn down and resignaled.

MBB procedures for manual and timer-based resignaling of the LSP, and for TE graceful shutdown, are supported.

The **tools>perform>router>mpls>update-path** command is not supported for mesh LSPs.

The **one-to-one** option under the **fast-reroute** command is also not supported.

If the TE database loses the router ID while the LSP is up, it will perform an update to the MPLS that states that the router ID is no longer in the TE database. This occurs whether the bypass backup path is activated or not. This will cause MPLS to tear down all mesh LSPs to this router ID. However, if the destination router is not a neighbor of the ingress LER and the user shuts down the IGP instance on the destination router, the router ID corresponding to the IGP instance will only be deleted from the TE database on the ingress LER after the LSA/LSP times out. If the user brings the IGP instance back up before the LSA/LSP times out, the ingress LER will delete and reinstall the same router ID at the receipt of the updated LSA/LSP. The RSVP-TE LSPs destined for this router ID will be deleted and re-established. All other failure conditions will cause the LSP to activate the bypass backup LSP or to go down without being deleted.

3.5.3.1 Multi-area and multi-instance support

A router that does not have TE links within a given IGP area or level will not have its router ID discovered in the TE database by other routers in this area or level. In other words, an auto-created LSP mesh cannot be signaled to a router that does not participate in the area or level of the ingress LER.

A mesh LSP can be signaled using TE links that belong to the same IGP area even if the router ID of the ingress and egress routers are interfaces reachable in a different area. In this case, the LSP is considered to be an intra-area LSP.

If multiple instances of IS-IS are configured on a router, each with its own router ID, the TE database on other routers will be able to discover TE links advertised by each instance. In this case, an instance of an LSP can be signaled to each router ID with a CSPF path computed using TE links within each instance.

If multiple instances of IS-IS are configured on a destination router, each with the same router ID, a single instance of LSP will be signaled from other routers. If the user shuts down one IGP instance, this will have no impact as long as the other IGP instances remain up. The LSP will remain up and will forward traffic using the same TE links. The same behavior exists with a provisioned LSP.

3.5.3.2 Mesh LSP name encoding

When the ingress LER signals the path of an auto-created mesh LSP, it includes the name of the LSP and the path name in the Session Name field of the Session Attribute object in the PATH message. The encoding is as follows:

Session Name: *<lspp-name::path-name>*, where the *lspp-name* component is encoded as follows:

TemplateName-DestIpv4Address-TunnelId

where **DestIpv4Address** is the address of the destination of the auto-created LSP.

3.5.4 Automatic creation of an RSVP-TE single-hop LSP

If the one-hop option is specified instead of a prefix policy, the **auto-lsp** command enables the automatic signaling of single-hop, point-to-point LSPs using the specified template to all directly connected neighbors. This LSP type is referred to as auto-created single-hop LSPs of type one-hop. Unlike the automatically created RSVP-TE LSP mesh, the automatically created single-hop RSVP-TE LSPs have no requirement for a prefix list to be referenced.

The first requirement is to create an LSP template containing the common parameters used to establish each single-hop LSP. The template must be created with the keyword **one-hop-p2p**:

config>router>mpls>lsp-template *template-name* one-hop-p2p

Upon creation of the template, CSPF is automatically enabled (and cannot be disabled), and the **hop-limit** is set to a value of two. The **hop-limit** defines the number of nodes the LSP may traverse, and since these are single-hop LSPs to adjacent neighbors, a limit of two is sufficient. The template must also reference a default path before it can be placed in the **no shutdown** state.

The next requirement is to trigger the creation of single-hop LSPs using the **auto-lsp lsp-template** command:

config>router>mpls>auto-lsp lsp-template *template-name* one-hop

The LSP and path parameters and options supported in an LSP template of type **one-hop-p2p** are the same as those in the LSP template of type **mesh-p2p**. The show command for **auto-lsp** will display the actual outgoing interface address in the "from" field.

The auto-created single-hop LSP can be signaled over both numbered and unnumbered RSVP-TE interfaces.

When the **one-hop** command is executed, the TE database keeps track of each TE link to a directly connected IGP neighbor whose router ID is discovered. MPLS then signals an LSP with a destination address matching the router ID of the neighbor and with a strict hop consisting of the address of the interface used by the TE link. The **auto-lsp** command with the **one-hop** option results in one or more LSPs signaled to the IGP neighbor.

Only the router ID of the first IGP instance of the neighbor that advertises a TE link causes the LSP to be signaled. If another IGP instance with a different router ID advertises the same TE link, no action is taken

and the existing LSP is kept up. If the router ID originally used disappears from the TE database, the LSP is kept up and is now associated with the other router ID.

The state of a single-hop LSP that is signaled displays the following behavior.

- If the interface used by the TE link goes down or BFD times out and the RSVP-TE interface is registered with BFD, the LSP path moves to the bypass backup LSP if the primary path is associated with one.
- If the association of the TE link with a router ID is removed from the TE database while the single-hop LSP is up, the single-hop LSP is torn down whether the bypass backup path is activated or not. This occurs if the interface used by the TE link is deleted or if the interface is shut down in the context of RSVP-TE.
- If the TE database loses the router ID while the LSP is up, it will perform two separate updates to MPLS, whether the bypass backup path is activated or not. The first one updates the loss of the TE link association, which will cause the single-hop LSP to be torn down. The other update states that the router ID is no longer in the TE database, which will cause MPLS to tear down all mesh LSPs to this router ID. A shutdown at the neighbor of the IGP instance that advertised the router ID will cause the router ID to be removed from the ingress LER node immediately after the last IGP adjacency is lost and not be subject to time-out as it is for a non-directly connected destination router.

All other feature behavior and limitations are the same as for an auto-created LSP mesh.

3.5.5 Automatic ABR selection for inter-area LSPs


Inter-area RSVP point-to-point LSPs support automatic area border router (ABR) selection at the ingress LER. The ABR does not need to be included as a loose hop in the LSP path definition.

CSPF can now compute all segments of a multi-segment, inter-area LSP path in one operation. Previously, MPLS made separate requests to CSPF for each segment.

For LSP path establishment, the explicit route object (ERO) in the PATH message is expanded on ABRs where the next hop is a loose hop in the LSP path definition. For ERO expansion to operate, the **cspf-on-loose-hop** command must be enabled under the **mpls** context on the ABR to allow the ABR to perform a CSPF calculation. If CSPF calculations are not performed, CSPF for the LSP path fails at the head-end node as TE information for links in another area are not available.

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

Figure 6: Automatic ABR selection for inter-area LSP



25052

CSPF for an inter-area LSP operates as follows:

1. CSPF in the ingress LER node determines that an LSP is inter-area by performing a route lookup with the destination address of a point-to-point LSP, such as 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 to be inter-area.
2. When the path of the LSP is empty, CSPF computes a single-segment, intra-area path to an ABR that advertised a prefix matching the destination address of the LSP.
3. If the path of the LSP contains one or more hops, CSPF computes a multi-segment, intra-area path including the hops that are in the area of the ingress LER node.
4. If all hops are in the area of the ingress LER, the calculated path ends on an ABR 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, the calculated path ends on an ABR that advertised a prefix matching the first-hop address that is not in the area of the ingress LER.
6. Note the following special case of a multi-segment, inter-area LSP. If CSPF hits a hop that can be reached via an intra-area path but that resides on an ABR, CSPF only calculates a path 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 on by having the ABR expand the remaining hops in the ERO.
7. If there is more than one ABR that advertised a prefix, CSPF calculates a path for all ABRs. Only the shortest path is withheld. If more than one path is the shortest path, CSPF picks a path randomly.

or based on the least-fill criterion if least-fill is enabled. If more than one ABR satisfies the least-fill criterion, CSPF also picks one path randomly.

8. The path for an intra-area LSP cannot exit and re-enter the local area of the ingress LER. This behavior was possible in prior implementations when the user specified a loose hop outside the local area or when the only available path was via TE links outside the local area.

3.5.5.1 Rerouting of inter-area LSPs

In prior implementations, an inter-area LSP path would have been rerouted if a failure or a topology change occurred in the local area or in a remote area while the ABR loose hop in the path definition was still up. If the transit/inter-area (exit) ABR failed or was put into node TE graceful shutdown, or if IS-IS went into overload mode, the LSP path would remain down at the ingress LER.

With automatic ABR selection, the ingress LER can reroute an inter-area LSP primary path via a different ABR in the following situations:

- When the local exit ABR fails, there are two cases to consider:
 - If the primary path is not protected at the ABR, and is therefore torn down by the previous hop in the path, then the ingress LER retries the LSP primary path via the ABR that currently has the best path for the destination prefix of the LSP.
 - If the primary path is protected at the ABR with a manual or dynamic bypass LSP, the ingress LER will receive a PathErr message with a notification of 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 MBB procedure to optimize the LSP primary path via the ABR that currently 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 procedure to optimize the LSP primary path via the ABR that currently 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 at the next timer or manual optimization of the LSP path if the IS-IS overload bit is received.

3.5.5.2 Behavior of MPLS options in inter-area LSPs

The automatic ABR selection for an inter-area LSP does not change the prior implementation of inter-area LSP behavior for many of the LSP-level and path-level options. However, there are a number of enhancements introduced by the automatic ABR selection feature:

- Features such as path bandwidth reservation and admin-groups continue to operate within the scope of all areas since they rely on propagating the parameter information in the PATH message across the area boundary.
- The TE graceful shutdown feature continues to support MBB of the LSP path to avoid the link or node that originated the PathErr message as long as the link or node is in the local area of the ingress LER. If the PathErr originated in a remote area, the ingress LER is not able to avoid the link or node when it performs the MBB since it computes the path to the local exit ABR only. However, there is an exception to this. An enhancement has been added to cause the upstream ABRs in the current path of the LSP to record the link or node to avoid and use the record in subsequent ERO expansions. This means that if the ingress LER computes a new MBB path that goes through the same exit ABR as the current path, and all ABRs upstream of the node or link that originated the PathErr message are also selected in the new MBB path when the ERO is expanded, then the new path will also avoid this link or node.

- MBB support has been expanded to avoid the ABR when the node is put into TE graceful shutdown.
- The **lsp>metric-type te** or **lsp-template>cspf use-te-metric** option in CSPF cannot be propagated across the area boundary and therefore operates within the scope of the local area of the ingress LER.
- The **srlg** option on the bypass LSP continues to operate locally at each PLR within each area. The PLR 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 with the automatic ABR selection feature.
- The **least-fill** option support with an inter-area LSP is introduced with the automatic ABR selection feature. When this option is enabled, CSPF applies the least-fill criterion to select the path segment to the exit ABR in the local area.
- The PLR 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 of the area, CSPF must return no path.
- With the automatic ABR selection feature, timer-based resignaling of the inter-area LSP path is supported and the path is resignaled if the cost of the path segment to the local exit ABR changes. The cost shown for the inter-area LSP at the ingress LER is the cost of the path segments to the ABR.

3.5.5.3 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 (OSPF backbone area). All prefixes in area 0 appear to be reachable via an intra-area path. However, the prefixes are not reachable since 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 of 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 since none of these links are directly connected to this node.

The inter-area LSP feature can identify when the destination of an LSP is reachable via a virtual link. In that case, CSPF automatically computes and signals an inter-area LSP via the ABRs 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.


3.5.6 ABR FRR protection for inter-area LSP

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

Manual bypass protection only works when a proper completely strict path is provisioned that avoids the ABR.

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

Figure 7: ABR protection using dynamic bypass LSP



25053

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

- The PLR must inspect the 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 route 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 message for the primary path) along with the constraint to exclude the protected ABR and the include/exclude admin groups of the primary path. If CSPF returns a path that can only go to an intermediate hop, then the PLR signals the dynamic bypass and automatically includes the XRO with the address of the protected ABR and propagates the admin-group constraints of the primary path into the Session Attribute object of the bypass LSP. Otherwise, the PLR signals the dynamic bypass directly to the merge point node with no XRO in the PATH message.
- If a node-protect dynamic bypass cannot be found or signaled, the PLR attempts a link-protect dynamic bypass LSP. As with the 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 only works if the RRO node-id also contains the interface address. Otherwise, the neighbor is not created once the bypass is activated by the PLR. The Path state then 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 7705 SAR, acting as a PLR, will not signal a one-to-one detour LSP for ABR protection. In addition, an ABR will reject a PATH message, received from a third party implementation, with a detour object and with the ERO having the next hop loose. This is performed whether the **cspf-on-loose** option is enabled or not on the 7705 SAR. In other words, the 7705 SAR, working as a transit ABR for the detour path, rejects the signaling of an inter-area detour backup LSP.

3.6 Preference option for standby secondary LSP paths

The **path-preference** command allows priority values to be assigned to standby secondary LSP paths. This command can only be used for secondary LSP paths that have been configured in standby mode.

When the primary LSP becomes inactive, the standby secondary LSP with the highest path priority (lowest **path-preference** value) is chosen from the qualifying standby secondary LSPs to become the active LSP. This functionality allows a user to choose a path for one of the standby secondary LSPs that may, for example, be over a more reliable link or over a link with a lower latency.

If multiple standby secondary LSP paths have the same priority value, the system selects the path with the lowest uptime.

3.7 RSVP-TE FRR

FRR is a mechanism to protect against RSVP-TE signaled LSP failures by reacting to these failures as soon as possible. FRR is set up from the iLER, which signals the transit routers to precompute their backup LSPs. FRR creates a precomputed backup LSP from each node in the LSP path. If a link or LSP between two routers fails, traffic is rerouted immediately onto the precomputed backup LSP.



Note: In order for FRR to work, CSPF must be enabled.

The 7705 SAR supports FRR facility backup and one-to-one backup.

Facility backup mode allows FRR to be enabled on an aggregate basis and protects a whole node or a whole link, regardless of the number of LSPs using that link. In other words, facility backup mode creates a common bypass tunnel to protect all LSP paths traversing a common facility path. It provides flexibility, faster provisioning, and faster convergence times compared with one-to-one backup or LSP redundancy. One-to-one backup allows FRR to be enabled on a per-LSP basis.

With both methods, MPLS switches build many possible detour routes on the nodes between the ingress and egress nodes of an LSP. The facility backup method creates a detour route between two nodes, called a bypass tunnel, which is a single tunnel that follows the primary LSP path except where the link or node has failed. Traffic then switches to the bypass tunnel. The bypass tunnel merges with the original LSP path at the merge point (MP) as soon as possible. The one-to-one backup method creates a detour route, called a detour LSP, for each LSP that needs to be rerouted. Unlike the bypass tunnel, the detour LSP takes the best path to the termination point, and does not merge with the original LSP as soon as possible. The detour LSPs of a one-to-one backup LSP can merge at a detour merge point (DMP), which can either be at the termination point or at a point along the primary LSP.

One of the major differences between facility and one-to-one backup is the scalability offered by the protection method. In facility backup mode, all LSPs of the same type are rerouted over the bypass tunnel. Hence they are all protected against the failure of a node or link in the network. In facility backup mode,

each LSR along the path verifies that it has a bypass tunnel available to meet its requirements; otherwise, if it can, it signals a new bypass tunnel based on the requirements. If a new LSP is configured for FRR facility backup, the existing backup tunnels are scanned and if any one of them can be used for recovery, it is preferred. If there are no common links, then a new bypass tunnel will be signaled, assuming that the LSP requirements can be met. One-to-one backup mode uses similar reroute and protection methods except a detour route is applied on a per-LSP basis.

The 7705 SAR uses CSPF to calculate the explicit route and dynamically signal the FRR LSP.

With facility backup mode, routers check the contents of the Record Route Object (RRO) in the received RESV message to determine the bypass tunnel endpoint in the FRR facility. For link protection, the router uses the RRO to check the IP address of the next-hop router attached to the far end of the link along with the label allocation information and to build the bypass tunnel. This label is preserved until the LSP is merged at the MP. For node protection, the router uses the RRO to determine the next-next-hop router and the label it is expecting. The collection of RRO information is enabled through the **record** and **record-label** options.

If, after this process, another LSP requests FRR using the facility backup method, the router checks and compares its session object to the existing session objects and if there is a match, the router binds that LSP to the same bypass tunnel. If there is no match, another bypass is created.

3.7.1 FRR terminology

The following table provides definitions of terms used for FRR.

Table 6: FRR terminology

Term	Definition
Backup path	The LSP that is responsible for backing up a protected LSP. A backup path can be a backup tunnel (facility backup) or a detour LSP (one-to-one backup).
Backup tunnel	The LSP that is used to back up one of the many LSPs in FRR facility (many-to-one) backup
Bypass tunnel	An LSP that is used to protect a set of LSPs passing over a common facility in FRR facility backup. A bypass tunnel can be configured manually or dynamically (see Dynamic and manual bypass LSPs).
CSPF	Constraint-based shortest path first
Detour route	Any alternate route that protects the primary path, such as a secondary path, FRR bypass tunnel, or FRR detour LSP. The term "detour route" should not be confused with the term "detour LSP". Detour route is a general term that refers to any alternate route, while detour LSP is a specific term that applies to one-to-one backup.
Detour LSP	The LSP that is used to reroute traffic around a failure in FRR one-to-one backup. The term "detour LSP" should not be confused with the term "detour route". Detour route is a general

Term	Definition
	term that refers to any alternate route, while detour LSP is a specific term that applies to one-to-one backup.
DMP	Detour merge point In the case of one-to-one backup, this is an LSR where multiple detours converge. Only one detour is signaled beyond that LSR.
Disjoint	See SRLG disjoint
Facility backup	A local repair method in which a single bypass tunnel is used to protect one or more LSPs that traverse the PLR, the resource being protected, and the Merge Point (in that order). Facility backup is distinct from a one-to-one backup tunnel, which has one backup path per protected path.
MP	Merge point The LSR where one or more backup tunnels rejoin the path of the protected LSP downstream of the potential failure. The same LSR may be both an MP and a PLR simultaneously.
NHOP bypass tunnel	Next-hop bypass tunnel A backup tunnel that bypasses a single link of the protected LSP
NNHOP bypass tunnel	Next-next-hop bypass tunnel A backup tunnel that bypasses a single node of the protected LSP
One-to-one backup	A local repair method in which a backup LSP is separately created for each protected LSP at a PLR
PLR	Point of local repair The head-end router of a backup tunnel or a detour LSP, where the term local repair refers to techniques used to repair an LSP tunnel quickly when a node or link along an LSP path fails
Primary path	An LSP that uses the routers specified by the path defined by the primary path-name command
Protected LSP	An LSP is protected at a given hop if it has one or more associated backup tunnels originating at that hop
Reroutable LSP	Any LSP for which the head-end router requests local protection
Secondary path	An LSP that protects a primary path that uses LSP redundancy protection rather than FRR protection
SRLG disjoint	A path is considered to be SRLG disjoint from a given link or node if the path does not use any links or nodes that belong to the same SRLG as the given link or node

3.7.2 Bypass resignal timer

When the bypass resignal timer is enabled, MPLS makes a request to CSPF for the best path for each dynamic bypass LSP originated on the node. The constraints of the first associated LSP primary path that originally triggered the signaling of the bypass LSP must be satisfied. In order to do this, MPLS saves the original Path State Block (PSB) of the LSP primary path, even if the path is torn down.

If CSPF returns no path or returns a new path that is equal in cost to the current path, the PSB associations are not updated. If CSPF returns a new path with a different cost from the current one, MPLS signals it.

When 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 whether the corresponding LSP primary path constraints are still satisfied by the new bypass LSP path. If the constraints are satisfied, the PSB association is moved to the new bypass LSP.

If the constraints are not satisfied, the PSB remains associated with the older bypass LSP and will be checked at the next background PSB re-evaluation or at the next timed or manual bypass reoptimization. Additionally, if the older bypass LSP is SRLG disjoint with a primary path that has the non-strict SRLG condition and the new bypass LSP is not SRLG disjoint, the PSB association is not moved.

If a 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) operation tears down all corresponding primary paths, which also causes the older bypass LSP to be torn down.

When the **bypass resignal timer** is configured, a PSB re-evaluation task is initiated that runs in the background of each RSVP-TE session to determine whether an existing manual or dynamic bypass is more optimal for that session. If the PSB re-evaluation task finds a more optimal bypass, it moves the PSB association to it. 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 will, however, audit 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, due to the PSB constraints not being met by the new signaled bypass LSP path, will be re-evaluated by the background task against all existing manual and dynamic bypass LSPs.

The background re-evaluation task also checks for PSBs that have requested a node-protect bypass LSP but are currently associated with a link-protect bypass LSP, as well as PSBs that have requested FRR protection and have no association. The background task is in addition to the attempt made when an RESV message is received on the protected LSP path, which ensures the association is completed faster.

This feature is not supported with inter-area dynamic bypass LSPs.

3.7.3 FRR behavior

The FRR MPLS facility backup method and one-to-one backup method are configured on the ingress LER (iLER) by using the **fast-reroute** command.

The behavior of an LSP at an iLER with both FRR and a standby LSP path configured is as follows.

- When a downstream detour route (alternative path) becomes active at a Point of Local Repair (PLR):
The iLER switches to the standby LSP path as soon as it is notified of the reroute. If the primary LSP path is subsequently repaired at the PLR, the LSP switches back to the primary path. If the standby

path goes down, the LSP is switched back to the primary path, even though the primary path is still on the detour route at the PLR.

- If the primary path goes down at the iLER while the LSP is on the standby path, the detour route at the iLER is torn down and, for one-to-one backup detour routes, a "path tear" is sent for the detour route. In other words, the detour route at the iLER does not protect the standby LSP. If and when the primary LSP is again successfully resigaled, the iLER detour route will be restarted.
- When the primary LSP fails at the iLER:
The LSP switches to the detour route. If the primary path undergoes a global revertive recovery, the LSP switches back to the primary path. If the LSP is on the detour route and the detour route fails, the LSP is switched to the standby path.
- Administrative groups are not taken into account when creating the detour routes for LSPs.

3.7.4 Dynamic and manual bypass LSPs

Users can disable dynamic bypass creation on a per-node basis using the **config>router>mpls>dynamic-bypass** command. Disabling dynamic bypass means that manual bypass is enabled. Dynamic bypass is enabled by default.

Dynamic bypass tunnels are implemented as per RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*. When an LSP is signaled and the Local Protection flag in the Session_attribute object is set, or the FRR object in the PATH message indicates that facility backup is desired, the PLR establishes a bypass tunnel to provide node and link protection. If there exists a bypass LSP that merges with the protected LSP at a downstream node, and if this LSP satisfies the constraints in the FRR object, then this bypass tunnel is selected and used. The **frr-object** command specifies whether facility backup is signaled in the FRR object.

The manual bypass feature allows an LSP to be preconfigured from a Point of Local Repair (PLR) that will be used exclusively for bypass protection. When a PATH message for a new LSP requests bypass protection, the node first checks for a manual bypass tunnel that satisfies the path constraints. If one is found, it is selected and used. If no manual bypass tunnel is found, the 7705 SAR dynamically signals a bypass LSP in the default behavior. To configure a manual bypass LSP, use the **bypass-only** option in the **config>router>mpls>lsp lsp-name [bypass-only]** command.


When a PLR activates a bypass backup LSP and subsequently receives a RESV refresh message for the original primary LSP path reservation over the restored interface, the PLR does not generate a ResvErr packet downstream. In addition, the MP node, once it becomes active, does not propagate a downstream ResvErr message received packet for the original primary LSP path reservation.

See [Configuring manual bypass tunnels](#) for configuration information.

3.7.4.1 Bypass LSP selection rules for the PLR

The following figure shows an example of a network used to illustrate the LSP selection rules for a PLR bypass scenario.

Figure 8: Bypass tunnel node example



20123

The PLR uses the following rules to select a bypass LSP from among multiple bypass LSPs (manually and dynamically created) when establishing the primary LSP path or when searching for a bypass for a protected LSP that does not have an association with a bypass tunnel:

1. The MPLS/RSVP-TE task in the PLR node checks for an existing manual bypass tunnel that satisfies the constraints. If the PATH message for the primary LSP path indicated that node protection is desired, which is the default LSP FRR setting at the head-end node, then the MPLS/RSVP-TE task searches for a node-protect bypass LSP. If the PATH message for the primary LSP path indicated that link protection is desired, then it searches for a link-protect bypass LSP.
2. If multiple manual bypass LSPs satisfying the path constraints exist, the PLR will prefer a manual bypass LSP terminating closer to the PLR over a manual bypass LSP terminating further away. If multiple manual bypass LSPs satisfying the path constraints terminate on the same downstream node, the PLR selects the one with the lowest IGP path cost, or if there is a tie, it picks the first one available.
3. If none of the manual bypass LSPs satisfy the constraints and dynamic bypass tunnels have not been disabled on the PLR node, then the MPLS/RSVP-TE task in the PLR node checks to determine if any of the already established dynamic bypass LSPs of the requested type satisfy the constraints.
4. If none of the dynamic bypass LSPs satisfy the constraints, the MPLS/RSVP-TE 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 required, and no manual bypass was found after step 1 or no dynamic bypass LSP was found after three attempts to perform step 3, the MPLS/RSVP-TE task will repeat steps 1 to 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 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 desired, 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 LSPs that requested FRR protection but are not currently associated with a bypass tunnel, the PLR node—upon reception of an RESV refresh message on the primary LSP path—repeats steps 1 to 7.

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

If the user configures a bypass tunnel on Node B ([Figure 8: Bypass tunnel node example](#)) and dynamic bypass tunnels have been disabled, LSPs that had been previously signaled and that were not associated with any manual bypass tunnel (for example, none existed) will be associated with the manual bypass tunnel, if it is suitable. The node checks for the availability of a suitable bypass tunnel for each of the outstanding LSPs every time an 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 that had been previously signaled over dynamic bypass tunnels will not automatically be switched to the manual bypass tunnel, even if the manual bypass tunnel is a more optimized path. The user must perform a make-before-break switchover at the head end of these LSPs. The make-before-break process is enabled using the **adaptive** option.

If the manual bypass tunnel goes into the down state on 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 the "protection available" flag in the next RESV refresh message for each of these LSPs. If it cannot find one, it will keep checking for one every time an RESV message is received for each of the remaining LSPs. When the manual bypass tunnel is back up, the LSPs that did not find a match are associated back with this tunnel and the protection available flag is set starting in the next RESV refresh message.

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

3.7.4.2 FRR node protection (facility backup)


The MPLS fast reroute (FRR) functionality enables PLRs to be aware of the lack of node protection and lets them regularly probe for a node bypass via the **node-protect** command.

When enabled, the **node-protect** command provides node protection for the specified LSP. If node protection cannot be provided, link protection is attempted. If link protection cannot be provided, no protection is provided. When disabled via the **no** form of the command, link protection is attempted, and if link protection cannot be provided, no protection is provided.

For example, assume the following for the LSP scenario in [Figure 9: FRR node-protection example](#).

1. LSP_1 is between PE_1 and PE_2 (via P1 and P2), and has CSPF, FRR facility backup, and FRR node protection enabled.
2. P1 protects P2 with bypass nodes P1 - P3 - P4 - PE_4 - PE_3.
3. If P4 fails, P1 tries to establish the bypass node three times.
4. When the bypass node creation fails (there is no bypass route), P1 will protect link P1-P2.
5. P1 protects the link to P2 through P1 - P5 - P2.
6. P4 returns online.

Figure 9: FRR node-protection example



20124

LSP_1 had requested node protection, but due to lack of an available path it could only obtain link protection. Therefore, every 60 s, the PLR for LSP_1 will search for a new path that may be able to provide node protection. When P4 is back online and such a path is available, a new bypass tunnel will be signaled and LSP_1 will be associated with this new bypass tunnel.

3.7.5 Admin group support on facility bypass backup LSPs

Admin group support on facility bypass backup LSPs provides for the inclusion of the LSP primary path admin-group constraints in the computation of an FRR facility bypass backup LSP to protect the primary LSP path. Admin group constraints are honored by all nodes in the LSP path both for primary and FRR backup LSPs.

This feature is supported on primary paths of an RSVP point-to-point LSP in both intra-area and inter-area TE where applicable.

This feature is not supported on one-to-one detour backup LSPs.

3.7.6 FRR over unnumbered interfaces

When the PLR is the ingress LER node and the outgoing interface of the bypass LSP is unnumbered, the user must assign a borrowed IP address to the interface that is different from the system interface; otherwise, the bypass LSP will not come up.

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

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

- the C-Type value of the RSVP_HOP object has changed

- the C-Type is the value for the IF_ID RSVP_HOP object (C-Type = 3) and it has not changed, but the IF_ID TLV is different
- the IPv4 Next/Previous Hop Address field in the RSVP_HOP object has changed, regardless of the C-Type value

This behavior at the PLR and MP nodes is the same for both link protection and node protection FRR.



Note: If node protection FRR is enabled but the MP does not support an unnumbered interface, the PATH message is rejected at the MP and the path is torn down.

See [RSVP-TE support for unnumbered interfaces](#) for information about unnumbered interfaces.

3.8 Shared risk link groups

A shared risk link group (SRLG) represents a set of interfaces (or links) that share the same risk of failing because they may be subjected to the same resource failures or defects. Two examples where the same risk of failure exists are fiber links that share the same conduit, and multiple wavelengths that share the same fiber.

SRLGs are supported by both LSP redundancy protection and FRR protection. SRLGs allow the user to prepare a detour route that is disjoint from the primary LSP path. See [Disjoint and non-disjoint paths](#).

The SRLG feature ensures that a primary and secondary LSP path, or a bypass tunnel or detour LSP path, do not share SRLGs. That is, they do not share the same sets of links that are considered to have a similar (or identical) chance of failure.

To use SRLGs, the user first creates an SRLG by assigning one or more routers to the SRLG. Then, the user links the SRLG to an MPLS interface and enables the SRLG feature on the LSP path. SRLGs cannot be assigned to the system interface.

3.8.1 SRLGs for secondary LSP paths

SRLGs for secondary LSP paths apply when LSP redundancy protection is used.

When setting up the secondary path, enable the **srlg** option on the secondary path to ensure that CSPF includes the SRLG constraint in its route calculation. To make an accurate computation, CSPF requires that the primary LSP be established and in the up state (because the head-end LER needs the most current explicit route object (ERO) for the primary path, and the most current ERO is built during primary path CSPF computation). The ERO includes the list of SRLGs.

At the establishment of a secondary path with the SRLG constraint, the MPLS/RSVP-TE task queries CSPF again, which provides the list of SRLGs to be avoided. CSPF prunes all links having interfaces that belong to the same SRLGs as the interfaces included in the ERO of the primary path. If CSPF finds an eligible path, the secondary path is set up. If CSPF does not find an eligible path, MPLS/RSVP-TE keeps retrying the requests to CSPF.

3.8.2 SRLGs for FRR LSP paths

When setting up the FRR bypass or detour LSP, enable the **srlg-frr** option on FRR to ensure that CSPF includes the SRLG constraint in its route calculation. CSPF prunes all links that are in the SRLG being

used by the primary LSP during the calculation of the FRR path. If one or more paths are found, CSPF sets up the FRR bypass or detour LSP based on the best cost and signals the FRR LSP.

If there is no path found based on the above calculation and the **srlg-frr** command has the **strict** option set, then the FRR LSP is not set up and the MPLS/RSVP-TE task keeps trying to set up a path. If the **strict** option is not set, then the FRR LSP is set up based on the other TE constraints (that is, excluding the SRLG constraint).

3.8.3 Disjoint and non-disjoint paths

A path is considered to be SRLG disjoint from a given link (or node) if the path does not use any links (or nodes) that belong to the same SRLG as the given link (or node). Eligible disjoint paths are found by CSPF when the SRLG constraint is included in the CSPF route calculation (referred to as the strict SRLG condition).

When LSP redundancy is used, the secondary LSP is always signaled with a strict SRLG condition.

When FRR is used, the FRR bypass or detour LSP may have a strict or non-strict SRLG condition. If the **strict** option is used with the **srlg-frr** command, then the bypass LSP must be on the list of eligible paths found by the CSPF calculation that included the SRLG constraint. If the **strict** option is not used, then it is possible for the bypass or detour LSP to be non-disjoint. The non-disjoint case is supported only if the SRLG is not strict.

At the PLR, if an FRR tunnel is needed to protect a primary LSP, the priority order for selecting that FRR tunnel is as follows:

1. Manual bypass disjoint
2. Manual bypass non-disjoint (eligible only if **srlg-frr** is non-strict)
3. Dynamic bypass disjoint
4. Dynamic bypass non-disjoint (eligible only if **srlg-frr** is non-strict)

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 considered in the CSPF calculation.

3.8.4 Enabling disjoint backup paths

A typical application of the SRLG feature is to provide automatic setup of secondary LSPs or FRR bypass or detour LSPs, in order to minimize the probability that they share the same failure risks with the primary LSP path.

Figure 10: Disjoint primary and secondary LSPs illustrates SRLG when LSP redundancy is used, where SRLG_1 contains the interfaces that define links A-B, B-C, and C-D. The primary path uses these links to connect node A to node D. In the event of a failure along the primary path, the secondary path cannot use any of the links in SRLG_1 and takes the path from node A to nodes E, F, G, H, J, and D.

Figure 11: Disjoint FRR bypass LSPs illustrates SRLG when FRR bypass is used, where SRLG_1 is the same as in **Figure 10: Disjoint primary and secondary LSPs**. Since FRR bypass is used, the following possible reroutes may occur, depending on where the failure occurs:

- if node B fails, the bypass is from node A to nodes E, F, G, H, and C
- if node C fails, the bypass is from node B to nodes F, G, H, J, and D

- if link C-D fails, the bypass is from node C to nodes H, J, and D

The SRLG feature is supported on OSPF and IS-IS interfaces for which RSVP-TE is enabled.

The following steps describe how to enable SRLG disjoint backup paths for LSP redundancy and FRR.

LSP redundancy for primary/secondary (standby) SRLG disjoint configuration

- Create an SRLG group (similar to creating an admin group).
- Link the SRLG group to MPLS interfaces.
- Configure primary and secondary LSP paths, and enable SRLG on the secondary LSP path. The SRLG secondary LSP paths will always perform a strict CSPF query.


The setting of the **srlg-frr** command is irrelevant in this case (see the [srlg-frr](#) command).

FRR bypass tunnel or detour LSP SRLG disjoint configuration

- Create an SRLG group (similar to creating an admin group).
- Link the SRLG group to MPLS interfaces.
- Enable the **strict** option on the **srlg-frr** command, which is a system-wide command that forces the CSPF calculation for every LSP path to take any configured SRLG memberships into account.
- Configure primary FRR (facility backup or one-to-one backup) LSP paths. Each PLR will create a bypass or detour LSP that will only avoid the SRLG memberships configured on the primary LSP path egress interface. For one-to-one backup, detour-detour merging is out of the control of the PLR. The PLR will not ensure that the FRR detour will be prohibited from merging with a colliding detour LSP. For facility backup, given that there are several bypass types to bind to, the priority rules shown in [Disjoint and non-disjoint paths](#) are used.


Manually configured bypasses that do not use CSPF are not considered as possible backup paths.

Figure 10: Disjoint primary and secondary LSPs



20482

Figure 11: Disjoint FRR bypass LSPs



20483

3.9 RSVP-TE graceful shutdown

RSVP-TE graceful shutdown provides a method to reroute transit LSPs in a bulk fashion away from a node prior to maintenance of that node. A PathErr message with the error code "Local Maintenance on TE Link required Flag" (if the affected network element is a link) or the error code "Local node maintenance required" (if the affected network element is the node) is sent before the links or node are taken out of service.

When an LER receives the message, it performs a make-before-break on the LSP path to move the LSPs away from the links/nodes whose IP addresses are indicated in the PathErr message and reroute them. Affected link/node resources are flagged in the TE database so that other routers will signal LSPs using the affected resources only as a last resort.

Graceful shutdown can be enabled on a per-interface basis or on all interfaces on the node if the whole node must be taken out of service.

3.10 RSVP-TE support for unnumbered interfaces

Unnumbered interfaces are point-to-point interfaces that are not explicitly configured with a dedicated IP address and subnet; instead, they borrow (or link to) an IP address from another interface on the system (the system IP address, another loopback interface, or any other numbered interface) and use it as the source IP address for packets originating from the interface. For more information about support for unnumbered interfaces, see the 7705 SAR Router Configuration Guide, *Unnumbered Interfaces*.

Unnumbered IP interfaces can be used via RSVP-TE for signaling traffic engineering (TE) LSPs.

Supporting RSVP-TE over unnumbered interfaces requires the ability to:

- carry TE information over unnumbered links in IS-IS-TE or OSPF-TE extensions
- specify unnumbered interfaces in RSVP-TE signaling

An unnumbered IP interface is identified uniquely on a router in the network by the tuple (router ID, ifindex). An LSR at each end of the link assigns a system-wide unique interface index to the unnumbered interface. IS-IS, OSPF, MPLS (RSVP-TE, LDP), and OAM use this tuple to advertise the link information, signal LSPs over the interface, or send and respond to an MPLS echo request message over an unnumbered interface.

The borrowed IP address for an unnumbered interface is configured using the following CLI command, with the default value set to the system interface address: **config>router>interface>unnumbered {ip-int-name | ip-address}**.



Note: The borrowed IP address is used exclusively as the source address for IP packets that originate from the interface. For FRR, this address must be configured to an address different from the system interface in order for the FRR bypass LSP to come up at the ingress LER. See [RSVP-TE FRR](#) for information on FRR.

To support unnumbered TE links in IS-IS, a new sub-TLV of the extended IS reachability TLV is added, which encodes the link local identifiers and link remote identifiers as defined in RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*.

To support unnumbered TE links in OSPF, a new sub-TLV of the Link TLV is added, which encodes the link local identifiers and link remote identifiers as defined in RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*.

To support unnumbered TE links in RSVP-TE, a new sub-object of the Explicit Route Object (ERO) is added to specify unnumbered links and a new sub-object of the Route Record Object (RRO) is added to record that the LSP traversed an unnumbered link, as per RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*. As well, a new IF_ID RSVP_HOP object with a C-Type of 3 is added as per section 8.1.1 of RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions*. The IPv4 Next/Previous Hop Address field in the object is set to the borrowed IP interface address.

The unnumbered IP interface address is advertised by IS-IS-TE or OSPF-TE, and CSPF can include it in the computation of a path for a point-to-point LSP. However, this feature does not 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 that advertised an unnumbered interface index is obtained from the TE database. Therefore, if traffic engineering is disabled in IS-IS or OSPF, a non-CSPF LSP that has its next hop over an unnumbered interface will not come up at the ingress LER because the router ID of the neighbor that has the next hop of the PATH message cannot be looked up. The LSP path will remain operationally down with the error "noRouteToDestination". If a PATH message is received at the LSR for which traffic engineering was disabled and the next hop for the LSP is over an unnumbered interface, a PathErr message is sent back to the ingress LER with the error code of 24 "Routing Problem" and an error value of 5 "No route available toward destination".

All MPLS (RSVP-TE and LDP) features supported for numbered IP interfaces are supported for unnumbered interfaces, with the following exceptions:

- configuration of a router ID with a value other than system interface
- signaling of an LSP with an ERO based on a loose or strict hop using an unnumbered TE link in the path hop definition
- signaling of a one-to-one detour LSP over an unnumbered interface

- RSVP Hello messages and all Hello-related capabilities, such as Graceful-Restart Helper
- RSVP refresh reduction on an unnumbered interface

The unnumbered interface feature also extends the support of LSP ping and LSP traceroute to point-to-point LSPs that have unnumbered TE links in their path.

3.11 PCEP support for RSVP-TE LSPs

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 7705 SAR operates as a path computation 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. See [MPLS and RSVP-TE command reference](#) for command descriptions. See the [PCEP support for RSVP-TE LSPs](#) section in the [PCEP](#) chapter for information about using these commands.

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

3.12 Segment routing with traffic engineering (SR-TE)

Segment routing adds the ability to perform shortest path routing and source routing using the concept of abstract segments to IS-IS and OSPF routing protocols. A segment can represent a local prefix of a node, a specific adjacency of the node (interface/next hop), a service context, or a specific explicit path over the network. For each segment, the IGP advertises an identifier referred to as a segment ID (SID).

When segment routing is used together with the MPLS data plane, the SID is a standard MPLS label. A router forwarding a packet using segment routing will therefore push one or more MPLS labels.

Segment routing using MPLS labels can be used in both shortest path routing applications (see the 7705 SAR Routing Protocols Guide for information) and in traffic engineering (TE) applications, as described in this section.

The following are the objectives and applications of segment routing:

- ability for a node to specify a unicast shortest-route or source-routed forwarding path with the same mechanism; IGP can be reused to minimize the number of control plane protocols
- ability to use IGP-based MPLS tunnels without the addition of any other signaling protocol

- ability to tunnel services from ingress PE to egress PE with or without an explicit path and without requiring forwarding plane or control plane state in intermediate nodes
- ability to use Layer 3 spoke SDP interfaces to support multicast for segment routing. See the 7705 SAR Routing Protocols Guide, "Multicast for Segment Routing".
- FRR: ability to expand coverage of basic LFA to any topology with the use of a source-routed backup path; precomputation and setup of backup path without additional signaling
- support for LFA policies with shared-risk constraints, admin-groups, and link/node protection
- support for SR-TE entropy labels
- support for TE that includes loose/strict options, distributed and centralized TE, path disjointness, ECMP awareness, and limited or no per-service state on midpoint and tail-end routers
- support for fine-grained flow steering and service chaining via a centralized stateful path computation element (PCE) such as the one provided by the Nokia Network Services Platform (NSP)

3.12.1 SR-TE support

The following MPLS commands and modes are supported:

- global [router] MPLS-level commands and modes:
interface, lsp, path, shutdown, sr-te-resignal-timer
- LSP-level commands and modes:
bfd, bgp-transport-tunnel, exclude, hop-limit, include, label-stack-reduction, local-sr-protection, metric, path-computation-method, primary, metric-type, retry-limit, retry-timer, shutdown, to, from, vprn-auto-bind
- primary path-level commands and modes:
bandwidth, bfd, exclude, hop-limit, include, shutdown
- secondary path-level commands and modes:
bandwidth, bfd, exclude, hop-limit, include, path-preference, shutdown, standby

The following MPLS commands and modes are not supported:

- global MPLS-level commands and modes not applicable to SR-TE LSPs (configuration is ignored):
admin-group-frr, auto-lsp, bypass-resignal-timer, cspf-on-loose-hop, dynamic-bypass, frr-object, hold-timer, least-fill-min-thd, least-fill-reoptim-thd, logger-event-bundling, lsp-template, srlg-frr, static-lsp, static-lsp-fast-retry
- LSP-level commands and modes not supported with SR-TE LSPs (configuration is blocked):
adaptive, adspec, fast-reroute, least-fill, propagate-admin-group, rsvp-resv-style
- LSP-level commands and modes not supported with SR-TE LSPs (configuration is ignored):
igp-shortcut
- primary path-level commands and modes not supported with SR-TE LSPs (configuration is blocked):
adaptive, record, record-label
- secondary path-level commands and modes not supported with SR-TE LSPs (configuration is blocked):
adaptive, record, record-label

The user can associate an empty path or a path with strict or loose explicit hops with the primary paths of the SR-TE LSP using the **hop**, **primary**, and **secondary** CLI commands.

A hop that corresponds to an adjacency SID must be identified with its far-end host IP address (next hop) on the subnet. If the local-end host IP address is provided, this hop is ignored because this router can have multiple adjacencies (next hops) on the same subnet.

A hop that corresponds to a node SID is identified by the prefix address.

Details of processing the user-configured path hops are provided in [SR-TE LSP instantiation](#).

3.12.2 SR-TE LSP instantiation

When an SR-TE LSP is configured on the router, its path can be computed by the router or by an external TE controller referred to as a PCE. This feature works with the Nokia stateful PCE that is part of the Network Services Platform (NSP).

The 7705 SAR supports three different modes of operation configurable on a per- SR-TE LSP basis.

- When the path of the LSP is computed by the router acting as a PCC, the LSP is referred to as PCC-initiated and PCC-controlled.

A PCC-initiated and controlled SR-TE LSP has the following characteristics:

- can contain strict or loose hops, or a combination of both
- supports both a basic hop-to-label translation and a full CSPF as a path computation method
- has the capability to report an SR-TE LSP to synchronize the LSP database of a stateful PCE server using the **pce-report** option, but the LSP path cannot be updated by the PCE. The control of the LSP is maintained by the PCC.

- 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 SR-TE LSP supports the passive stateful mode, which enables the **path-computation-method pce** option for the SR-TE LSP so that the PCE can perform path computation at the request of the PCC only. The PCC retains control.

The capability exists to report an SR-TE LSP to synchronize the LSP database of a stateful PCE server using the **pce-report** option.

- When the path of the LSP is computed and 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 SR-TE LSP allows active stateful mode, which enables the **pce-control** option for the SR-TE LSP so that the PCE can perform path computation and updates following a network event without the explicit request from the PCC. The PCC delegates full control.

The user can configure the path computation requests only (PCE-computed) or both path computation requests and path updates (PCE-controlled) to the PCE for a specific LSP using the **path-computation-method pce** and **pce-control** commands.

The **path-computation-method pce** option sends the path computation request to the PCE instead of the local CSPF. When this option is enabled, the PCE acts in passive stateful mode for this LSP. In other words, the PCE can perform path computations for the LSP only at the request of the router. This 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 value is **no path-computation-method**.

The user can also enable the router's full local CSPF path computation method. See [SR-TE LSP path computation using local CSPF](#) for more details.

The **pce-control** option allows the router to delegate full control of the LSP to the PCE (PCE-controlled). Enabling this option means that the PCE is acting in active stateful mode for this LSP and the PCE can reroute the path following a failure or to reoptimize the path and update the router without requiring a request from the router.



Note:

- The user can delegate LSP path computation to either the local CSPF or can use an explicit path with hop-to-label translation.
- The user can delegate LSPs that have the LSP path computation to the PCE via the **path-computation-method pce** command. The LSP maintains its latest active path computed by the PCE or the router at the time it was delegated. The PCE will only make an update to the path at the next network event or reoptimization. The default value is **no pce-control**.

In all cases, the PCC LSP database is synchronized with the PCE LSP database using the PCEP path computation state report (PCRpt) message for LSPs that have the **pce-report** command enabled.

The global MPLS- level **pce-report** command can be used to enable or disable PCE reporting for all SR-TE LSPs for the purpose of LSP database synchronization. This configuration is inherited by all LSPs of a particular type. The PCE reports both CSPF and non-CSPF LSPs. The default value is disabled (**no pce-report**). This default value controls the introduction of the PCE into an existing network and allows the operator to decide if all LSP types need to be reported.

The LSP-level **pce-report** command overrides the global configuration for PCE reporting for an LSP. The default value is to inherit the global MPLS-level value. The **inherit** value reconfigures the LSP to inherit the global configuration for that LSP type.



Note: If PCE reporting is disabled for the LSP, either due to inheritance or due to LSP-level configuration, enabling the **pce-control** option for the LSP has no effect. To help troubleshoot this situation, operational values of both the **pce-report** and **pce-control** are added to the output of the LSP **show** commands.

For more information about configuring PCC-initiated and PCC-controlled LSPs, see [Configuring PCC-controlled, PCE-computed, and PCE-controlled SR-TE LSPs](#).

3.12.2.1 PCC-initiated and PCC-controlled LSPs

For PCC-initiated and PCC-controlled LSPs, the user configures the LSP name, primary path name, and optional secondary path name with the path information in the referenced path name, entering a full or partial explicit path with all or some hops to the destination of the LSP. Each hop is specified as an address of a node or an address of the next hop of a TE link. Optionally, each hop can be specified as a SID value that corresponds to the MPLS label to use on a hop. If this option is used, the whole path must consist of SIDs.

The router supports both full local CSPF and basic hop-to-label translation path computation methods for an SR-TE LSP. In addition, the user can configure a path for the SR-TE LSP by explicitly entering SID label values.

To configure the primary path or secondary path to always use a specific link whenever it is up, the strict hop must be entered as an address corresponding to the next hop of an adjacency SID. If the strict hop

corresponds to a loopback address, it is translated to an adjacency SID as explained below and therefore there is no guarantee that the same TE link is picked.

To use an SR-TE path that consists of unprotected adjacency SIDs, each hop of the path must be configured as a strict hop with the address matching the next hop of the adjacency SID and protection on each of these adjacencies must be disabled as explained in [SR-TE LSP path computation](#).

MPLS assigns a tunnel ID to the SR-TE LSP and a path ID to each new instantiation of the primary path, as for an RSVP-TE LSP. These IDs represent the MBB path of the same SR-TE LSP, which must coexist during the update of the primary path.



Note: The concept of MBB is not exactly accurate in the context of an SR-TE LSP because there is no signaling involved and therefore the new path information immediately overrides the older one.

When local CSPF is not configured, the router retains full control of the path of the LSP and the full or partially explicit path is instantiated as is and no other constraint (such as SRLG, admin-group, hop-count, or bandwidth) is checked. Only the LSP path label stack size is checked by MPLS against the maximum value configured for the LSP after the TE database (TE-DB) hop-to-label translation returns the label stack. See [SR-TE LSP path computation](#) for more information about this check.

The ingress LER performs the following steps to resolve the user-entered path before programming it in the data path:

1. MPLS passes the path information to the TE-DB, which converts the list of hops into a label stack by scanning the TE-DB for adjacency and node SID information that belongs to the router or link identified by each hop address. If the conversion is successful, the TE-DB will return the actual selected hop SIDs plus labels as well as the configured path hop addresses that were used as the input for this conversion.

Details of this step are as follows:

- A loose hop with an address matching any interface (loopback or not) of a router (identified by router ID) is always translated to a node SID. If the prefix matching the hop address has a node SID in the TE-DB, it will be selected by preference. If not, the node SID of any loopback interface of the same router that owns the hop address is selected. In the latter case, the lowest IP address of that router that has a /32 prefix-SID is selected.
- A strict hop with an address matching any interface (loopback or not) of a router (identified by router ID) is always translated to an adjacency SID. If the hop address matches the host address reachable in a local subnet from the previous hop, the adjacency SID of that adjacency is selected. If the hop address matches a loopback interface, it is translated to the adjacency SID of any link from the previous hop that terminates on the router owning the loopback. The adjacency SID label of the selected link is used.

In both cases, it is possible to have multiple matching previous hops if the interface is a LAN interface. If there are multiple hops, the adjacency SID with the lowest interface address is selected.

- All IGP instances are scanned from the lowest to the highest instance ID, beginning with IS-IS instances and then the OSPF instance; not only the IGP instance that resolved the prefix of the destination address of the LSP in the RTM is used. For the first instance where all specified path hop addresses can be translated, the label stack is selected. The hop-to-SID/label translation tool does not support paths that cross area boundaries. All SID/labels of a given path are therefore taken from the same IGP area and instance.



Note: For the hop-to-label translation to operate, the user must enable TE on the network links by adding the network interfaces to MPLS and RSVP. In addition, the user must

enable the **traffic-engineering** option on all participating router IGP instances. If a router has the **database-export** option enabled in the participating IGP instances to populate the TE-DB with the learned IGP link-state information, then enabling of the **traffic-engineering** option is not required. For consistency, it is recommended that the **traffic-engineering** option always be enabled.

2. The ingress LER validates the first hop of the path to determine the outgoing interface and next hop to forward the packet to, and programs the data path according to the following conditions.
 - If the first hop corresponds to an adjacency SID (host address of next hop on the link's subnet), the adjacency SID label is not pushed. In other words, the ingress LER treats forwarding to a local interface as a push of an implicit null label.
 - If the first hop is a node SID of a downstream router, the node SID label is pushed.

In both cases, the SR-TE LSP tracks and uses the SR shortest path tunnel of the SID of the first hop.
3. If the router is configured as a PCC and has a PCEP session to a PCE, the router sends a PCRpt message to update the PCE with the Up state and the RRO object for each LSP that has the **pce-report** option enabled. The PE router does not set the delegation control flag to keep LSP control. The state of the LSP is now synchronized between the router and the PCE.

3.12.2.1.1 Guidelines for using PCC-initiated and PCC-controlled LSPs

The 7705 SAR supports both a full local CSPF and a basic hop-to-label translation path computation methods for a SR-TE LSP. In addition, the user can configure a path for the SR-TE LSP by explicitly entering SID label values.

The ingress LER can use any of the following methods to detect if a path is down, or is not optimal, and to take corrective action:

- Failure of the top SID detected via a local failure or an IGP network event. In this case, the LSP path goes down and MPLS retries it.
- Timeout of the seamless BFD session when enabled on the LSP and the failure-action is set to the value of **failover-or-down**. In this case, the path goes down and MPLS retries it.
- Receipt of an IGP link event in the TE database. In this case, MPLS performs an as-needed reoptimization of the paths of all SR-TE LSPs if the user enabled the MPLS level command **sr-te-resignal resignal-on-igp-event**. This capability only works when the path computation method is the local CSPF. It allows the ingress LER not only to detect a single remote failure event which causes packets to drop but also a network event which causes a node SID to reroute and therefore forwarding packets on a potentially sub-optimal TE path.
- Performing a manual or timer based resignal of the SR-TE LSP. This applies only when the path computation method is the local CSPF. In this case, MPLS reoptimizes the paths of all SR-TE LSPs.

With both the hop-to-label path computation method and the user configured SID labels, the ingress LER does not monitor network events that affect the reachability of the adjacency SID or node SID used in the label stack of the LSP, except for the top SID. As a result, the label stack may not be updated to reflect changes in the path except when seamless BFD is used to detect failure of the path. It is therefore recommended to use this type of SR-TE LSP in the following configurations only:

- empty path
- path with a single node-SID loose-hop

- path of an LSP to a directly-connected router (single-hop LSP) with an adjacency-SID or a node-SID loose/strict hop
- strict path with hops consisting of adjacencies explicitly configured in the path as IP addresses or SID labels.

3.12.2.2 PCC-initiated and PCE-computed/controlled LSPs

In the PCC-initiated and PCE-computed/controlled LSP mode of operation, the ingress LER uses PCEP to communicate with a PCE-based external TE controller (also referred to as the PCE). The router instantiates a PCEP session to the PCE. The router is referred to as the path computation client (PCC).

When the user enables the **path-computation-method pce** option for one or more SR-TE LSPs, the PCE performs path computations at the request of the PCC, which is referred to as passive stateful mode. If the user enables the **pce-control** option for an LSP, the PCE can also perform both path computation and periodic reoptimization of the LSP path without an explicit request from the PCC. This is referred to as active stateful mode.

For the PCC to communicate with a PCE about the management of the path of an SR-TE LSP, the router implements the extensions to PCEP in support of segment routing (see [PCEP](#) for more information). This feature works with the Nokia stateful PCE, which is part of the NSP.

The following steps describe configuring a PCC-initiated SR-TE LSP when passive or active control is given to the PCE:

1. The SR-TE LSP configuration is created on the PE router using the CLI or NSP NFM-P.
The configuration dictates which PCE stateful mode is desired: active (**pce-control** option enabled) or passive (**path-computation-method pce** enabled and **pce-control** disabled).
2. The 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. The PCC on the router tracks the association of {PLSP-ID, SRP-ID} to {tunnel-ID, path-ID} and uses the latter to communicate with MPLS about a specific path of the LSP.
3. The PE router does not validate the entered path. While the PCC can include the IRO objects for any loose or strict hop in the configured LSP path in the path computation request (PCReq) message to the PCE, the PCE ignores the IRO objects and computes the path with the other constraints.
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. It also includes the LSP object with the assigned PLSP-ID. At this point, the PCC does not delegate control of the LSP to the PCE.
5. The PCE computes a new path, reserves the bandwidth, and returns the path in a path computation reply (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 the computed metric value if any, and the BANDWIDTH object.



Note: In order for 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 verify that the traffic-engineering option is enabled in the IGP. This causes the 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 CSM and the data path with the new path.

Up to this step, the PCC and PCE are using passive stateful PCE procedures. The next steps synchronize the LSP database of the PCC and PCE for both PCE-computed and PCE-controlled LSPs. They 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 the Up state and the RRO as confirmation, including the LSP object with the unique PLSP-ID. For a PCE-controlled LSP, the PE router also sets a 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 reoptimization, the PCE computes a new path for a PCE-controlled LSP and returns it in a path computation update (PCUpd) message with the new ERO. It includes the LSP object with the same unique PLSP-ID assigned by the PCC and 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 CSM and the data path with the new path.
10. The PE router sends a new PCRpt message to update the PCE with the Up state and the RRO as confirmation. The state of the LSP is now synchronized between the router and the PCE.
11. If the user makes any configuration change to the PCE-computed or PCE-controlled LSP, MPLS requests the PCC to revoke delegation in a PCRpt message (PCE-controlled only), and then MPLS and the PCC follow the above steps to convey the changed constraint to the PCE, which will result in a new path programmed into the data path, the LSP databases of the PCC and PCE to be synchronized, and the delegation to be returned to the PCE.

For SR-TE LSPs, MBB is not supported. Therefore, the PCC first tears down the LSP and sends a PCRpt message to the PCE with the remove flag set to 1 before following this configuration change procedure.



Note: The above procedures are followed when the user performs a **no shutdown** on a PCE-controlled or PCE-computed LSP. The starting point is an administratively down LSP with no active paths.

The following steps are for an LSP with an active path.

- If the user enabled the **path-computation-method pce** option on a PCC-controlled LSP that has an active path, no action is performed until the next time the router needs a path for the LSP following a network event or an LSP parameter change. At that point, the procedures above are followed.
- If the user enabled the **pce-control** option on a PCC-controlled or PCE-computed LSP that has an active path, the PCC will issue a PCRpt message to the PCE with the Up state and the RRO of the active path. The PCC will set the delegation control flag to delegate control to the PCE. The PCE will keep the active path of the LSP and will not update until the next network event or reoptimization. At that point, the procedures above are followed.

The PCE supports the computation of disjoint paths for two different LSPs originating or terminating on the same or different PE routers. To indicate this constraint to the PCE, the user must configure the PCE path profile ID and path group ID that the LSP belongs to. These parameters are passed transparently by the PCC to the PCE and are therefore opaque data to the router. The user can configure the path profile and path group using the **path-profile profile-id [path-group group-id]** command.

The association of the optional path group ID is to allow the PCE to determine which profile ID this path group ID must be used with. One path group ID is allowed per profile ID. The user can, however, enter the same path group ID with multiple profile IDs by executing this command multiple times. A maximum of five entries of **path-profile [path-group]** can be associated with the same LSP. More details of the operation of the PCE path profile are provided in the [PCEP](#) chapter.

3.12.3 SR-TE LSP path computation

For PCC-controlled SR-TE LSPs, CSPF is supported on the router. See [SR-TE LSP path computation using local CSPF](#) for details about the full local CSPF path computation method. By default, the path is computed using the hop-to-label translation method. In this case, MPLS makes a request to the TE-DB to get the label corresponding to each hop entered by the user in the primary path of the SR-TE LSP. See [PCC-initiated and PCC-controlled LSPs](#) for details of the hop-to-label translation.

The user can configure the path computation request of a CSPF-enabled SR-TE LSP to be forwarded to a PCE instead of the local router CSPF by enabling the **path-computation-method pce** option, as described in [SR-TE LSP instantiation](#). The user can further delegate the reoptimization of the LSP to the PCE by enabling the **pce-control** option. In both cases, the PCE is responsible for determining the label required for each returned explicit hop and includes this in the SR-ERO.

In all cases, the user can configure the maximum number of labels that the ingress LER can push for a particular SR-TE LSP by using the **max-sr-labels** command.

This command is used to set a limit on the maximum label stack size of the SR-TE LSP primary path to allow room to insert additional transport, service, and other labels when packets are forwarded in a particular context.

CLI syntax:

```
config>router>mpls>lsp>max-sr-labels label-stack-size [additional-frr-labels labels]
```

The **max-sr-labels label-stack-size** value should be set to account for the desired maximum label stack of the primary path of the SR-TE LSP. Its range is 1 to 11 and the default value is 6.

The value in **additional-frr-labels labels** should be set to account for additional labels inserted by remote LFA or Topology Independent LFA (TI-LFA) for the backup next hop of the SR-TE LSP. Its range is 0 to 4 labels with a default value of 1.

The sum of both label values represents the worst-case transport of SR label stack size for this SR-TE LSP and is populated by MPLS in the TTM such that services and shortcut applications can check it to decide if a service can be bound or a route can be resolved to this SR-TE LSP. More details of the label stack size check and requirements in various services and shortcut applications are provided in [Service and shortcut application SR-TE label stack check](#).

The maximum label stack supported by the router is discussed in [Data path support](#). The maximum label stack is always signaled by the PCC in the PCEP Open object as part of the SR-PCE-CAPABILITY TLV. It is referred to as the Maximum Stack Depth (MSD).

In addition, the per-LSP value for the **max-sr-labels label-stack-size** option, if configured, is signaled by the PCC to the PCE in the SID Depth value in a METRIC object for both a PCE-computed LSP and a PCE-controlled LSP. The PCE will compute and provide the full explicit path with TE links specified. If there is no path with the number of hops lower than the MSD value or the SID Depth value (if signaled), a reply with no path will be returned to the PCC.

For a PCC-controlled LSP, if the label stack returned by the TE-DB hop-to-label translation exceeds the per-LSP maximum SR label stack size, the LSP is brought down.

3.12.3.1 Service and shortcut application SR-TE label stack check

Each service and shortcut application on the router performs a check of the resulting net label stack after pushing all the labels required for forwarding the packet in that context. The MPLS module populates each

SR-TE LSP in the TTM with the maximum transport label stack size, which consists of the sum of the values in **max-sr-labels** *label-stack-size* and **additional-frr-labels** *labels*.

Each service or shortcut application then adds the additional, context-specific labels, such as service label and NGE label, required to forward the packet in that context, and checks that the resulting net label stack size does not exceed the maximum label stack supported by the router.

If the check succeeds, the service is bound or the prefix is resolved to the SR-TE LSP. If the check fails, the service will not bind to this SR-TE LSP. Instead, the service will either find another SR-TE LSP or another tunnel of a different type to bind to, if the user configured the use of other tunnel types. Otherwise, the service will go down.

When the service uses an SDP with one or more SR-TE LSPs (up to eight), the spoke SDP bound to this SDP will remain operationally down as long as at least one SR-TE LSP fails the check. In this case, the spoke SDP flag "labelStackLimitExceeded" will be displayed in the **show** output of the service. As well, the prefix will not get resolved to the SR-TE LSP and will either be resolved to another SR-TE LSP or another tunnel type or become unresolved.

The value of **additional-frr-labels** *labels* is checked against the maximum value across all IGP instances of the parameter **frr-overhead**. The **frr-overhead** parameter value is computed within an IGP instance as shown in the following table. For more information about FRR overhead, see the *7705 SAR Routing Protocols Guide*, "Segment Routing in Shortest Path Forwarding".

Table 7: Parameter values for *frr-overhead*

Condition	Parameter value
segment-routing is disabled in the IGP instance	0
segment-routing is enabled but remote-lfa is disabled and ti-lfa is disabled	0
segment-routing is enabled and remote-lfa is enabled but ti-lfa is disabled	1
segment-routing is enabled and ti-lfa is enabled, regardless of whether remote-lfa is enabled or disabled	ti-lfa max-sr-frr-labels <i>label</i>

When the user configures or changes the configuration of **additional-frr-labels**, MPLS ensures that the new value accommodates the **frr-overhead** parameter value across all IGP instances.

For example:

- The user configures the **config>router>isis>loopfree-alternates>remote-lfa** command.
- The user creates a new SR-TE LSP or changes the configuration of an existing SR-TE LSP as follows: **mpls>lsp>max-sr-labels 10 additional-frr-labels 0**.
- Performing a **no shutdown** of the new LSP or changing the existing LSP configuration will be blocked because the IS-IS instance enabled remote LFA, which requires one additional label on top of the 10 SR labels of the primary path of the SR-TE LSP.

If the check is successful, MPLS adds **max-sr-labels** and **additional-frr-labels** and checks that the sum is lower than or equal to the maximum label stack supported by the router. MPLS then populates the value of {**max-sr-labels** + **additional-frr-labels**}, along with tunnel information in the TTM, and also passes **max-sr-labels** to the PCEP module.

Conversely, if the user tries a configuration change that results in a change to the computed **frr-overhead**, the IGP will check that all SR-TE LSPs can properly account for the overhead; otherwise, the change is rejected. On the IGP, enabling **remote-lfa** may cause the **frr-overhead** value to change.

For example:

- An MPLS LSP is administratively enabled and has **mpls>lsp>max-sr-labels 10 additional-frr-overhead 0** configured.
- The current configuration in IS-IS or OSPFv2 has the **loopfree-alternates** command disabled.
- The user attempts to configure **loopfree-alternates>remote-lfa** for IS-IS or OSPFv2. This changes **frr-overhead** to 1.

This configuration change would be blocked.

When the user configures the **ti-lfa** command, the **max-sr-frr-labels value** parameter is used to limit the search for the LFA backup next hop, as follows:

- 0 – the IGP LFA SPF restricts the search to the TI-LFA backup next hop that does not require a repair tunnel, meaning that the P node and Q node are the same and match a neighbor. This is also the case when both P and Q nodes match the advertising router for a prefix. For information about P nodes and Q nodes, see *draft-francois-rtgwg-segment-routing-ti-lfa-04 (Topology Independent Fast Reroute using Segment Routing)*.
- 1 to 3 – the IGP LFA SPF widens the search to include a repair tunnel to a P node that is connected to the Q nodes with zero to two hops for a total of three labels maximum: one node SID-to-P node and two adjacency SIDs from the P node to the Q node. If the P node is a neighbor of the computing node, its node SID is compressed, meaning that up to three adjacency SIDs can separate the P and Q nodes.
- 2 (default) – this corresponds to a repair tunnel to a non-adjacent P node that is adjacent to the Q node. If the P node is a neighbor of the computing node, the node SID of the P node is compressed and the default value of two labels corresponds to two adjacency SIDs between the P and Q nodes.

When the user attempts to change the **max-sr-frr-labels** parameter to a value that results in a change to the computed FRR overhead, the IGP checks that all SR-TE LSPs can properly account for the overhead based on the configuration of the LSP **max-sr-labels** and **additional-frr-labels** values; otherwise, the change is rejected.

The FRR overhead is computed by the IGP and its value is shown in [Table 7: Parameter values for frr-overhead](#).

The above LFA commands allow the user to enable the base LFA feature with the **loopfree-alternates** command, and to add remote LFA with the **remote-lfa** command and TI-LFA with the **ti-lfa** command. For more information, see the *7705 SAR Routing Protocols Guide*, "Segment Routing in Shortest Path Forwarding".

3.12.4 SR-TE LSP path computation using local CSPF

This section describes full local CSPF path computation for SR-TE LSP paths.

By default, SR-TE LSP paths use the hop-to-label translation method. The **path-computation-method** command overrides the default and allows the path-computation method to be set to local CSPF or PCE. The PCE path computation method is not supported with the SR-TE LSP template.

3.12.4.1 Extending MPLS and TE database CSPF support to SR-TE LSPs

The following MPLS and TE database features extend CSPF support to SR-TE LSPs:

- IPv4 SR-TE LSP
- local CSPF on both primary and secondary standby paths of an IPv4 SR-TE LSP
- support path computation in single area OSPFv2 and IS-IS IGP instances
- computes full explicit TE paths using TE links as hops and returning a list of adjacency SIDs. The details of the CSPF path computation are provided in [SR-TE specific TE-DB changes](#).
- use random path selection in the presence of ECMP paths that satisfy the LSP and path constraints. Least-fill path selection is not required.
- provide an option to reduce or compress the label stack such that the adjacency SIDs corresponding to a segment of the explicit path are replaced with a node SID whenever the constraints of the path are met by all the ECMP paths to that node SID. For more details, see [SR-TE LSP path label stack reduction](#).
- use legacy TE link attributes as in RSVP-TE LSP CSPF
- Uses timer reoptimization of all paths of the SR-TE LSP that are in the operational up state. This differs from the RSVP-TE LSP resignal timer feature that reoptimizes the active path of the LSP only. MPLS provides the current path of the SR-TE LSP and TE-DB updates the total IGP or TE metric of the path, checking the validity of the hops and labels as per current TE-DB link information. CSPF then calculates a new path and provides both the new and metric updated current path back to MPLS. MPLS programs the new path only if the total metric of the new computed path is different than the updated metric of the current path, or if one or more hops or labels of the current path are invalid. Otherwise, the current path is considered one of the most optimal ECMP paths and is not updated in the data path. Timer resignal applies only to the CSPF computation method and not to the ip-to-label computation method.
- use manual reoptimization of a path of the SR-TE LSP. In this case, the new computed path is always programmed even if the metric or SID list is the same.
- Supports ad hoc (as needed) reoptimization. As-needed resignaling of all SR-TE LSPs is triggered if one or more IGP link down events are received in the TE-DB.

After the reoptimization is triggered, the behavior is the same as the timer-based resignal or the delay option of the manual resignal. MPLS forces the expiry of the resignal timer and asks the TE-DB to reevaluate the active paths of all SR-TE LSPs. The re-evaluation consists of updating the total IGP or TE metric of the current path, checking the validity of the hops and labels, and computing a new CSPF for each SR-TE LSP. MPLS programs the new path only if the total metric of the new computed path is different from the updated metric of the current path, or if one or more hops or labels of the current path are invalid. Otherwise, the current path is considered one of the most optimal ECMP paths and is not updated in the data path.

- support unnumbered interfaces in the path computation. There is no support for configuring an unnumbered interface as a hop in the path of the LSP. The path can be empty or include hops with the address of a system or loopback interface, but path computation can return a path that uses TE links corresponding to unnumbered interfaces.
- support admin-group, hop-count, IGP metric, and TE-metric constraints
- Bandwidth constraint is not supported because SR-TE LSP does not have an LSR state to book bandwidth. The **bandwidth** parameter, when enabled on the LSP path, has no impact on local CSPF path calculation. However, the **bandwidth** parameter is passed to PCE when it is the selected path computation method. PCE reserves bandwidth for the SR-TE LSP path accordingly.

3.12.4.2 SR-TE specific TE-DB changes

When the **traffic-engineering** command is enabled in an OSPFv2 instance, only local and remote TE-enabled links are added into the TE-DB. A TE-link is a link that has one or more TE attributes added to it in the MPLS interface context. Link TE attributes are TE metric, bandwidth, and membership in an SRLG or an admin group.

To allow local CSPF to compute an SR-TE LSP with SR-enabled links that do not have TE attributes, the 7705 SAR implements the following changes:

- OSPFv2 passes all links, whether they are TE-enabled or SR-enabled, to the TE-DB, as currently performed by IS-IS.
- TE-DB relaxes the link back-check when performing a CSPF calculation to ensure that there is at least one link from the remote router to the local router. Because OSPFv2 advertises the remote link IP address or remote link identifier only when a link is TE-enabled, the strict check about the reverse direction of a TE-link cannot be performed if the link is SR-enabled but not TE-enabled.

As a result, even if an interface is administratively shut down in MPLS, an SR-TE LSP path that uses this interface does not go operationally down.

3.12.4.3 SR-TE LSP specific CSPF changes

The local CSPF for an SR-TE LSP is performed in two phases. Phase 1 computes a fully explicit path with all TE links to the destination specified, as in the case of an RSVP-TE LSP. If the user has enabled label stack reduction or compression for this LSP, Phase 2 is applied to reduce the label stack so that adjacency SIDs corresponding to a segment of the explicit path are replaced with a node SID whenever the constraints of the path are met by all the ECMP paths to that node SID. For more details, see [SR-TE LSP path label stack reduction](#).

The CSPF computation algorithm for the fully explicit path in Phase 1 remains mostly unchanged from its behavior in RSVP-TE LSP.

The meaning of a strict and loose hop in the path of the LSP is the same as in CSPF for RSVP-TE LSP. A strict hop means that the path from the previous hop must be a direct link. A loose hop means the path from the previous hop can traverse intermediate routers.

A loose hop may be represented by a set of back-to-back adjacency SIDs if not all paths to the node SID of that loose hop satisfy the path TE constraints. This is different from the IP-to-label path computation method where a loose hop always matches a node SID because no TE constraints are checked in the path to that loose hop.

When the label stack of the path is reduced or compressed, a strict hop may be represented by a node SID if all the links from the previous hop satisfy the path TE constraints. This is different from the IP-to-label path computation method where a strict hop always matches an adjacency SID.

The first phase of CSPF returns a full explicit path with each TE link specified all the way to the destination. The label stack may contain protected or unprotected adjacency SIDs. The user can configure the type of adjacency protection for the SR-TE LSP using a CLI command as described in [SR-TE LSP path protection](#).

The 7705 SAR does not support the origination of a global adjacency SID. If received from a third-party router implementation, it is added into the TE database but is not used in any CSPF path computation.

3.12.4.3.1 SR-TE LSP path protection

SR-TE LSP allows the user to configure whether the path of the LSP must use protected or unprotected adjacencies exclusively for all links of the path.

When the 7705 SAR forms an IGP adjacency over a link and the **segment-routing** context is enabled in the IGP instance, the static or dynamic label assigned to the adjacency is advertised in the link adjacency SID sub-TLV. By default, an adjacency is always eligible for LFA/RLFA/TI-LFA protection and the B-flag in the sub-TLV is set. The presence of a B-flag does not reflect the instant state of the availability of the adjacency LFA backup; it reflects that the adjacency is eligible for protection. The SR-TE LSP using the adjacency in its path still comes up if the adjacency does not have a backup programmed in the data path at that instant. Use the **config>router>isis>interface>no sid-protection** command to disable protection. When protection is disabled, the B-flag is cleared and the adjacency is not eligible for protection by LFA, remote LFA, or TI-LFA.

The local CSPF on the 7705 SAR can use all local and remote SIDs to compute a path for an SR-TE LSP based on the needed local protection property.

The 7705 SAR applies the following path protection behaviors when local CSPF is enabled on an SR-TE LSP:

- If the **local-sr-protection** command is not enabled or is set to **preferred**, the local CSPF prefers a protected adjacency over an unprotected adjacency whenever both exist for a TE link. This is done on a link-by-link basis after the path is computed based on the LSP path constraints. This means that the protection state of the adjacency is not used as a constraint in the path computation. It is only used to select a SID among multiple SIDs after the path is selected. The computed path can combine both types of adjacencies.

If multiples ECMP paths satisfy the constraints of the LSP path, one path is selected randomly and the SID selection above applies. There is no check if the selected path has the highest number of protected adjacencies.

- If the **local-sr-protection** command is set to **mandatory**, CSPF uses it as an additional path constraint and selects protected adjacencies exclusively in computing the path of the SR-TE LSP.

If no path satisfies the other LSP path constraints and consists of all TE links with protected adjacencies, the path computation returns no path.

- If the **local-sr-protection** command is set to **none**, CSPF uses it as an additional path constraint and selects unprotected adjacencies exclusively in computing the path of the SR-TE LSP.

If no path satisfies the other LSP path constraints and consists of all TE links with unprotected adjacencies, the path computation returns no path.

The **local-sr-protection** command impacts PCE-computed and PCE-controlled SR-TE LSPs. When the **local-sr-protection** command is set to the default value **preferred** or to the explicit value of **mandatory**, the local-protection-desired flag (L-flag) in the LSPA object in the PCReq (Request) message or in the PCRpt (Report) message is set to a value of 1.

When the **local-sr-protection** command is set to **none**, the local-protection-desired flag (L-flag) in the LSPA object is cleared. The PCE path computation checks this flag to decide if protected adjacencies are used in preference to unprotected adjacencies (L-flag set) or must not be used at all (L-flag clear) in the computation of the SR-TE LSP path.

3.12.4.3.2 SR-TE LSP path label stack reduction

Label stack reduction is used to reduce the label stack so ingress PE routers with a lower maximum SID depth (MSD) can still work. It can also be used to spray packets over ECMP paths to an intermediate node SID when all these paths satisfy the constraints of the SR-TE LSP path. This feature can be useful even if the resulting label stack is not reduced.

When label stack reduction is enabled on an LSP, local CSPF performs a second phase where the label stack reduction algorithm is used to replace adjacency SIDs corresponding to a segment of the explicit path with a node SID whenever the constraints of the path are met by all the ECMP paths to that node SID.

The label stack reduction algorithm uses the following procedure:

1. In Phase 1, the CSPF returns up to three fully explicit ECMP paths that are eligible for label stack reduction. These paths are equal cost from the point of view of the IGP metric or TE metric as configured for that SR-TE LSP.
2. Each fully explicit path of the SR-TE LSP that is computed in Phase 1 of the CSPF is split into a number of segments that are delimited by the user-configured loose or strict hops in the path of the LSP. Label stack reduction is applied to each segment separately.
3. Label stack reduction in Phase 2 consists of traversing the CSPF tree for each ECMP path returned in Phase 1, then attempting to find the farthest node SID in a path segment that can be used to summarize the entire path up to that node SID. This requires that all links of ECMP paths are able to reach the node SID from the current node on the CSPF tree to satisfy all the TE constraints of the SR-TE LSP paths. In this case, ECMP is based on the IGP metric because this is what routers use in the datapath when forwarding a packet to the node SID.

If the TE metric is enabled for the SR-TE LSP, one of the constraints is that the TE metric must be the same value for all the IGP metric ECMP paths to the node SID.

4. CSPF in Phase 2 selects the first candidate ECMP path from Phase 1 that reduces the label stack and also complies with the maximum label stack size constraint configured by the **sr-te-cspf** command.
5. The CSPF path computation in Phase 1 always avoids a loop over the same hop and the label stack reduction algorithm prevents a path from looping over the same hop because of the normal routing process. For example, it checks if the same node is involved in the ECMP paths of more than one segment of the LSP path and builds the label stack to avoid this situation.
6. During the MBB procedure of a timer or the manual reoptimization of an SR-TE LSP path, the TE-DB performs the following steps in addition to the initial path computation:
 - MPLS provides the TE-DB with the current working path of the SR-TE LSP.
 - The TE-DB updates the path's metric based on the IGP or TE link metric (if the TE metric enabled for the SR-TE LSP).
 - For each adjacency SID, the TE-DB verifies that the related link and SID are still in its database and that the link fulfills the LSP constraints. If so, it picks up the current metric.
 - For each node SID, the TE-DB verifies that the related prefix and SID are still available, and if so, checks that all the links on the shortest IGP path to the node owning the node SID fulfill the SR-TE LSP path constraints. This step reuses the same checks detailed in step 3.
 - CSPF computes a new path with or without label stack reduction as described in steps 1, 2, and 3.
 - The TE-DB returns both paths to MPLS. MPLS always programs the new path in the case of a manual reoptimization. MPLS compares the metric of the new path to the current path and if different, programs the new path in the case of a timer reoptimization.

7. The TE-DB sends the reduced path ERO and label stack to MPLS, along with the following information:
 - a list of SRLGs of each hop in the ERO, represented by a node SID, including the SRLGs used by links in all ECMP paths to reach that node SID from the previous hop
 - the cost of each hop in the ERO represented by an adjacency SID. The cost corresponds to the IGP metric or TE metric (if the TE metric is enabled for the SR-TE LSP) of that link.
 - the cost of each hop in the ERO represented by a node SID, which corresponds to the cumulated IGP metric or TE metric (if the TE metric is enabled for the SR-TE LSP) to reach the node SID from the previous hop using the fully explicit path computed in Phase 1.
 - the total cost or computed metric of the SR-TE LSP path. This consists of the cumulated IGP metric or TE metric (if TE metric enabled for the SR-TE LSP) of all hops of the fully explicit path computed in Phase 1 of the CSPF.
8. If label stack reduction is disabled, the values of the **max-sr-labels** and the **hop-limit** are applied to the full explicit path in Phase 1. The minimum of the two values is used as a constraint in the full explicit path computation. If the resulting ECMP paths net hop-count in Phase 1 exceeds this minimum value, the TE-DB does not return a path to MPLS.
9. If label stack reduction is enabled, the values of the **max-sr-labels** and the **hop-limit** are both ignored in Phase 1. However, the **max-sr-labels** value is used as a constraint in Phase 2. If the Phase 2 reduction of all candidate paths results in a net label stack size that exceeds the **max-sr-labels** value, the TE-DB does not return a path to MPLS.
10. The label stack reduction algorithm uses a node SID to replace a segment of the SR-TE LSP path; using a prefix SID with the N-flag clear is not supported.

3.12.4.3.3 Interaction with SR-TE LSP path protection


Label stack reduction is only attempted when the path protection **local-sr-protection** command is disabled or configured as **preferred**.

If **local-sr-protection** is configured as **none** or **mandatory**, the **label-stack-reduction** command is ignored, and the fully explicit path computed in Phase 1 is returned by the TE-DB CSPF routine to MPLS. This is because the node SID that is used to replace an adjacency SID can be unprotected or protected by LFA, based on local configuration on each router that resolves the node SID. However, this information is not directly advertised into the TE-DB. Therefore, CSPF cannot enforce the protection constraint requested along the path to that node SID.

3.12.4.3.4 Examples of SR-TE LSP path label stack reduction

The following figure shows a metro aggregation network with three levels of rings for aggregating regional traffic from edge ring routers to a PE router.

Figure 12: Label stack reduction in a 3-tier ring topology



The path of the highlighted LSP uses admin groups to force the traffic eastwards or westwards over the 3-ring topologies such that it uses the longest path possible. Assume all links in the bottom-most ring1 have admin-group=east1 for the eastward direction and admin-group=west1 for the westward direction.


Similarly, links in the middle ring2 have admin-group=east2 and admin-group=west2, and links in the top-most ring3 have admin-group=east3 and admin-group=west3. To achieve the longest path shown, the LSP or path should have an include statement: include east1 west2 east3. The fully explicit path computed in Phase 1 of CSPF results in label stack of size 18.

The label stack reduction algorithm searches for the farthest node SID in that path, which can replace a segment of the strict path while maintaining the admin-group constraints. The reduced label stack contains the SID adjacency MSR1-MSR2, the found node SIDs plus the node SID of the destination for a total of four labels to be pushed on the packet (the label for the adjacency MSR1-MSR2 is not pushed):

{N-SID MNR2, N-SID of MNR3, N-SID of MJR8, N-SID of PE1}

The following figure shows a topology example that creates two TE planes by applying a common admin group to all links of a specified plane. There are a total of four ECMP paths to reach PE2 from PE1: two within the red plane and two within the blue plane.

Figure 13: Label stack reduction in the presence of ECMP paths



For an SR-TE LSP from PE1 to PE2, which includes the red admin group as a constraint, Phase 1 of CSPF results in two fully explicit paths using adjacency SID of the red TE links:

path 1 = {PE1-P1, P1-P2, P2-P3, P3-PE2}

path 2 = {PE1-P1, P1-P4, P4-P3, P3-PE2}

Phase 2 of CSPF finds node SID of P3 as the farthest hop it can reach directly from PE1 while still satisfying the constraint to include the red admin-group constraint. If the node SID of PE2 is used as the only SID, then traffic would also be sent over the blue links. The reduced label stack is:

{P3 Node-SID=300, PE2 Node-SID=20}.

The resulting SR-TE LSP path combines the two explicit paths from Phase 1 into a single path with ECMP support.

3.12.5 SR-TE LSP paths using explicit SIDs

The 7705 SAR allows an explicit SID value to be configured for each hop in an SR-TE LSP path. The SID value for an SR-TE LSP hop is configured using the **sid-label** parameter under the **config>router>mpls>path>hop** context.

The **sid-value** parameter specifies an MPLS label value for that hop in the path.

The 7705 SAR supports a path computation method per LSP while allowing a mix-and-match configuration for primary and secondary paths. For example, primary and secondary paths can be different types: IP-to-hop paths or SID-based paths.

A path containing SID label hops is used even if the **path-computation-method** is configured for the LSP. The path computation method configured at the LSP level is ignored when explicit SIDs are used in

the path. This means that the router can bring up a path containing SID hops even if the LSP has path computation enabled.

When SIDs are explicitly configured for a path, the user must provide all of the necessary SIDs to reach the destination. The router does not validate whether the whole label stack provided is correct other than checking that the top SID is programmed locally. A path can come up even if it contains SIDs that are invalid. The user or controller programming the path should ensure that the SIDs are correct.

A path must consist of either all SIDs or all IP address hops.

Paths containing explicit SID values can only be used by SR-TE LSPs.

3.12.6 SR-TE LSP protection

The router supports local protection of a particular segment of an SR-TE LSP and end-to-end protection of the complete SR-TE LSP.

Each path is locally protected along the network using LFA or remote-LFA next hop whenever possible. The protection of a node SID reuses the LFA and remote LFA features introduced with segment routing shortest path tunnels; the protection of an adjacency SID has been added to the 7705 SAR in the specific context of an SR-TE LSP to augment the protection level. The user must enable the **loopfree-alternates>remote-lfa** command in IS-IS or OSPF.

An SR-TE LSP has state at the ingress LER only. The LSR has state for the node SID and adjacency SID, whose labels are programmed in the label stack of the received packet and which represent the part of the ERO of the SR-TE LSP on this router and downstream of this router. In order to provide protection for an SR-TE LSP, each LSR node must attempt to program a link-protect or node-protect LFA next hop in the ILM record of a node SID or an adjacency SID, and the LER node must do the same in the LTN record of the SR-TE LSP. The following are details of the behavior.

- If the ILM record is for a node SID of a downstream router that is not directly connected, the ILM of this node SID points to the backup NHLFE computed by the LFA SPF and programmed by the SR module for this node SID. Depending on the topology and LFA policy used, this can be a link-protect or node-protect LFA next hop.

This behavior is already supported in the SR shortest path tunnel feature at both the LER and LSR. Therefore, an SR-TE LSP that transits at an LSR and that matches the ILM of a downstream node SID automatically takes advantage of this protection when enabled. If required, node SID protection can be disabled under the IGP instance by excluding the prefix of the node SID from the LFA.

- If the ILM is for a node SID of a directly connected router, the LFA SPF only provides link protection. The ILM or LTN record of this node SID points to the backup NHLFE of this LFA next hop. An SR-TE LSP that transits at an LSR and that matches the ILM of a neighboring node SID automatically takes advantage of this protection when enabled.



Note: Only link protection is possible in this case because packets matching this ILM record can either terminate on the neighboring router owning the node SID or can be forwarded to different next hops of the neighboring router, that is, to different next next-hops of the LSR providing the protection. The LSR providing the connection does not have context to distinguish among all possible SR-TE LSPs and therefore can only protect the link to the neighboring router.

- If the ILM or LTN record is for an adjacency SID, it is treated as in the case of a node SID of a directly connected router.

When protecting an adjacency SID, the PLR first tries to select a parallel link to the node SID of the directly connected neighbor. That is the case when the node SID is reachable over parallel links. The selection is based on lowest interface ID. If no parallel links exist, regular LFA/remote LFA algorithms are applied to find a loopfree path to reach the node SID of the neighbor via other neighbors.

The ILM or LTN for the adjacency SID must point to this backup NHLFE and will benefit from FRR link protection. As a result, an SR-TE LSP that transits at an LSR and matches the ILM of a local adjacency SID automatically takes advantage of this protection when enabled.

- At the ingress LER, the LTN record points to the SR-TE LSP NHLFE, which points to the NHLFE of the SR shortest path tunnel to the node SID or adjacency SID of the first hop in the ERO of the SR-TE LSP. The FRR link or node protection at the ingress LER is inherited directly from the SR shortest path tunnel.

If an adjacency to a neighbor fails, the IGP withdraws the advertisement of the link TLV information as well as its adjacency SID sub-TLV. However, the LTN or ILM record of the adjacency SID must be kept in the data path for a sufficient period of time to allow the ingress LER to compute a new path after the IGP converges. If the adjacency is restored before the timer expires, the timer is aborted as soon as the new ILM or LTN records are updated with the new primary and backup NHLFE information. By default, the ILM/LTN and NHLFE information is kept for a period of 15 s.

The adjacency SID hold timer is configured using the **adj-sid-hold** command and activated when the adjacency to the neighbor fails due to the following conditions:

- the network IP interface went down due to a link or port failure or due to the user performing a shutdown of the port
- the user shuts down the network IP interface in the **config>router** or **config>router>ospf/isis** context

The adjacency SID hold timer is not activated if the user deletes an interface in the **config>router>ospf/isis** context.



Note:

- The adjacency SID hold timer does not apply to the ILM or LTN of a node SID, because NHLFE information is updated in the data path as soon as the IGP is converged locally and a new primary and LFA backup next hops have been computed.
- The label information of the primary path of the adjacency SID is maintained and reprogrammed if the adjacency is restored before the timer expires. However, the backup NHLFE may change when a new LFA SPF is run while the adjacency ILM is being held by the timer running. An update to the backup NHLFE is performed immediately following the LFA SPF and may cause packets to drop.
- A new protect group ID (PG-ID) is assigned each time an adjacency comes back up. This PG-ID is used by the ILM of the adjacency SID and the ILMs of all downstream node SIDs that resolve to the same next hop.

While protection is enabled globally for all node SIDs and local adjacency SIDs when the user enables the **loopfree-alternates** command in IS-IS or OSPF at the LER and LSR, there are applications where the user wants traffic to never divert from the strict hop computed by CSPF for an SR-TE LSP. In that case, the user can disable protection for all adjacency SIDs formed over a particular network IP interface using the **sid-protection** command.

The protection state of an adjacency SID is advertised in the B-FLAG of the IS-IS or OSPF Adjacency SID sub-TLV. No mechanism exists in PCEP for the PCC to signal to the PCE the constraint to use only adjacency SIDs, which are not protected. The path profile ID is configured in the PCE with the no-protection constraint.

3.12.6.1 Local protection

Each path can be locally protected through the network using LFA, remote LFA, or TI-LFA next hop whenever possible. The protection of a node SID reuses the LFA and remote LFA features introduced with SR shortest path tunnels. The protection of an adjacency SID has been added to the 7705 SAR in the specific context of an SR-TE LSP to augment the protection level. The user must enable the LFA, remote LFA, or TI-LFA option in IS-IS or OSPF.

This behavior is supported in the SR shortest path tunnel feature at both the LER and LSR. An SR-TE LSP that transits at an LSR and that matches the ILM of a downstream node SID automatically takes advantage of this protection when enabled. If required, node SID protection can be disabled under the IGP instance by excluding the prefix of the SID node from LFA.

3.12.6.2 End-to-end protection

This section provides a brief introduction to end-to-end protection for SR-TE LSPs. See [Seamless BFD for SR-TE LSPs](#) for detailed information about protection switching using seamless BFD (S-BFD) and a configured failure action.

The 7705 SAR supports end-to-end protection for SR-TE LSPs using secondary or standby paths. Standby paths (secondary paths with standby enabled) are permanently programmed in the data path, while secondary paths are only programmed when they are activated. S-BFD is used for end-to-end connectivity checking. The **failure-action failover-or-down** command under the **bfd** context of the LSP configures a switchover from the currently active path to an available standby or secondary path if the S-BFD session fails on the currently active path. If S-BFD is not configured, the router that is local to a segment can only detect failures of the top SID for that segment.

End-to-end protection with S-BFD can be combined with local protection but it is recommended that the S-BFD control packet timers be set to a value high enough to allow sufficient time for any local protection action for a particular segment to complete without triggering S-BFD to go down on the end-to-end LSP path.

To prevent failures between the paths of an SR-TE LSP (for example, to prevent a failure of a primary path that affects its standby backup path), disjoint paths should be configured.



Note: The 7705 SAR does not support shared risk link groups (SRLGs) for SR-TE. If the active/standby paths are not explicitly hopped throughout the network, they may merge into a single point of failure.

3.12.7 Seamless BFD for SR-TE LSPs

The 7705 SAR supports seamless BFD (S-BFD). Unlike LSP BFD, S-BFD does not rely on the traditional BFD session bootstrapping process or session state at the tail end of a session. Instead, when S-BFD is initialized the system selects a set of discriminators for the reflector or initiator function.

One S-BFD reflector is configured per system in the **config>bfd>seamless-bfd** context. A mapping between reflector discriminators and their IP addresses is configured on the initiator in the **config>router>bfd>seamless-bfd** context.

For information about seamless BFD, see the 7705 SAR Router Configuration Guide, *Seamless BFD*.

This section describes the application of S-BFD to SR-TE LSPs and the LSP configuration required for this feature.

S-BFD is supported in the following SR-TE contexts:

- PCC-initiated:
 - SR-TE LSP level
 - SR-TE primary path
 - SR-TE secondary and standby paths

3.12.7.1 Configuration of S-BFD on SR-TE LSPs

For PCC-initiated or PCC-controlled LSPs, the head end of an S-BFD session is configured under the SR-TE LSP context, the SR-TE primary path context, or the SR-TE secondary path context by using the following commands:

- **config>router>mpls>lsp *lsp-name* sr-te>bfd**
- **config>router>mpls>lsp *lsp-name* sr-te>primary>bfd**
- **config>router>mpls>lsp *lsp-name* sr-te>secondary>bfd**

The remote discriminator value is determined by passing the **to** address of the LSP to BFD, which then matches it to a mapping table of peer IP addresses to reflector remote discriminators. If there is no match for the **to** address of the LSP, a BFD session is not established on the LSP or path.



Note: A remote peer IP address-to-discriminator mapping must exist prior to bringing an LSP administratively up.

The referenced BFD template must specify parameters that are consistent with an S-BFD session; for example, the endpoint type must be **np**.

S-BFD can be configured at the LSP level, as follows:

```
configure>router>mpls>lsp <name> sr-te
  bfd
    [no] bfd-enable
    [no] bfd-template
    [no] failure-action
    [no] wait-for-up-timer <seconds>
  exit
```

When S-BFD is configured at the LSP level, separate S-BFD sessions with the same configuration are enabled on all primary, secondary, and standby paths of the LSP.

Alternatively, S-BFD can be configured on the primary path or secondary path of the LSP, as follows:

```
configure>router>mpls>lsp <name> sr-te
  primary <name>
    bfd
      [no] bfd-enable
      [no] return-path-label <label-value>
      [no] bfd-template <name>
      [no] wait-for-up-timer <seconds>
    exit
```

```
configure>router>mpls>lsp <name> sr-te
  secondary <name>
    bfd
      [no] bfd-enable
```

```
[no] return-path-label <label-value>
[no] bfd-template <name>
[no] wait-for-up-timer <seconds>
exit
standby
```

If desired, the S-BFD can be configured with the return path specified using the **[no] return-path-label/label-value** command.

The **wait-for-up-timer** command only applies if the configured failure action is **failover-or-down**. For more information, see [Support for BFD failure action with SR-TE LSPs](#).

3.12.7.2 Support for BFD failure action with SR-TE LSPs

S-BFD provides a mechanism to check the data path forwarding for an SR-TE LSP. If an S-BFD session fails, the LSP can be brought operationally down when the **failure-action** command is configured with the **failover-or-down** option. When the **failure-action** command is not configured, an S-BFD failure will only raise a trap. The **failure-action** command is available in the following context:

```
configure>router>mpls
  lsp <name> sr-te
    bfd
      failure-action failover-or-down
      no failure-action
```

The **failure-action** command is configured at the LSP level. It can be configured whether S-BFD is applied at the LSP level or the individual path level. The **failure-action** command can be configured even if the BFD template is not yet configured.

For LSPs configured with a primary path and a secondary or standby path and a failure action of **failover-or-down**, the following points apply:

- The path is held in an operationally down state when its S-BFD session is down.
- If all paths are operationally down, the SR-TE LSP is taken operationally down and a trap is generated.
- If S-BFD is enabled at the LSP or active path level, a switchover from the active path to an available path is triggered if the S-BFD session fails on the active path (primary or standby).
- If S-BFD is not enabled on the active path and this path is shut down, a switchover is triggered.
- If S-BFD is enabled on the candidate standby or secondary path, this path is only selected if S-BFD is up.
- An inactive standby path configured with S-BFD is only available to become active if it is not operationally down; that is, its S-BFD session is up and all other criteria for it to become operational are true. The path is held in an inactive state if the S-BFD session is down.
- The system does not revert to the primary path or start a reversion timer when the primary path is either administratively down or operationally down, because the S-BFD session is not up or down for any other reason.

For LSPs configured with only one path and a failure action of **failover-or-down**, the following points apply:

- The path is held in an operationally down state when its S-BFD session is down.
- If the path is operationally down, the LSP is taken operationally down and a trap is generated.



Note: S-BFD packets can still be sent on an operationally down SR-TE LSP.

For LSPs configured with one or more paths but no configured failure action, a BFD trap is raised when the LSP goes down. The path state does not change.

3.12.7.2.1 SR-TE LSP state changes and failure actions based on S-BFD

When a path is first configured with S-BFD, it is held operationally down until BFD comes up (subject to the BFD wait time).

The BFD **wait-for-up-timer** is started when BFD is first enabled on a path or when an existing S-BFD session transitions from up to down. If this timer expires before S-BFD is up, the path is torn down and the LSP retry timer is started.

In the S-BFD up-to-down case, if there is only one path, the LSP is torn down when S-BFD fails and then deprogrammed when the **wait-for-up-timer** expires.

If all the paths of an LSP are operationally down due to S-BFD, the LSP is taken operationally down and the BFD **wait-for-up-timer** is started for each path. If one or more paths do not have S-BFD configured or are otherwise not down, the LSP is not taken operationally down.

When an existing S-BFD session fails on a path and the failure action is **failover-or-down**, the configured failure action is activated, the path is put into the operationally down state, and a trap is raised. The state and reason code are displayed with the **show>router>bfd>seamless-bfd** command.

3.12.7.3 S-BFD operational considerations

A minimum control packet timer transmit interval of 10 ms can be configured. To maximize the reliability of S-BFD connectivity checking in scaled scenarios with short timers, situations where BFD may go down due to normal changes of the next hop of an LSP path at the head end must be avoided. It is therefore recommended that LFA not be configured at the head-end LER when using S-BFD with sub-second timers. When LFA is not configured, protection of the SR-TE LSP is still provided end-to-end by the combination of S-BFD connectivity checking and primary or secondary path protection.

Similar to LDP and RSVP functionality, S-BFD uses a single path for a loose hop; multiple S-BFD sessions for each of the ECMP paths or spraying of S-BFD packets across the paths is not supported. S-BFD is not down until all the ECMP paths of the loose hop go down.



Note: With very short control packet timer values in scaled scenarios, S-BFD may bounce if the next hop that the path is currently using goes down. This is because it takes time for BFD to be updated to use another next hop in the ECMP set.

3.12.8 Static route resolution using SR-TE LSPs

Static route packets can be forwarded to an indirect next hop over an SR-TE LSP programmed in the TTM with the following static route tunnel binding command:

CLI syntax:

```
config>router>static-route-entry ip-prefix/prefix-length [mcast]
      indirect ip-address
      tunnel-next-hop
```

```

        resolution {any | disabled | filter}
        resolution-filter
        [no] sr-te
            [no] lsp lsp-name
        exit
    exit
exit

```

3.12.9 BGP label route resolution using SR-TE LSPs

An SR-TE LSP programmed in the TTM can be used for resolving the next hop of a BGP IPv4 label route with the following BGP transport tunnel command:

CLI syntax:

```

config>router>bgp>next-hop-res>
label-route-transport-tunnel
    [no] family {label-ipv4 | vpn}
        resolution {any | disabled | filter}
        resolution-filter
        [no] sr-te
    exit
exit

```

3.12.10 Service packet forwarding using SR-TE LSPs

An SDP sub-type of the MPLS encapsulation type allows service binding of up to eight SR-TE LSPs programmed in the TTM by MPLS. The following example shows how to bind an SR-TE LSP to an MPLS SDP:

Example:

```

configure service sdp 100 mpls create
config>service>sdp$ sr-te-lsp lsp-name

```

The destination address of all LSPs must match the SDP far-end address. Service data packets are sprayed over the set of LSPs in the SDP using the same procedures as for tunnel selection in ECMP. In all cases, the SDP can only spray packets over a maximum of eight next hops. Each SR-TE LSP can, however, have up to eight next hops at the ingress LER when the first segment is a node SID-based SR tunnel. The SDP selects one next hop from each SR-TE LSP until the maximum number of eight next hops for the SDP is reached.

The **tunnel-far-end** option is not supported. In addition, the **mixed-lsp-mode** option does not support the **sr-te** tunnel type.

The signaling protocol for the service labels for an SDP using an SR-TE LSP can be configured to static (**off**), T-LDP (**tldp**), or BGP (**bgp**).

An SR-TE LSP can be used in VPRN autobind with the following commands:

CLI syntax:

```

config>service>vprn>
auto-bind-tunnel
    resolution {any | disabled | filter}
    resolution-filter

```

```

        [no] sr-te
    exit
exit

```

Both VPN-IPv4 and VPN-IPv6 (6VPE) are supported in a VPRN service using segment routing transport tunnels with the **auto-bind-tunnel** command.

This **auto-bind-tunnel** command is also supported with BGP EVPN service, as shown below:

CLI syntax:

```

config>service>vpls>bgp-evpn>mpls>
    auto-bind-tunnel
        resolution {any | disabled | filter}
        resolution-filter
            [no] sr-te
        exit
    exit
exit

```

The following service contexts are supported with SR-TE LSPs:

- VLL, VPLS, IES/VPRN spoke SDP and interface, and r-VPLS
- Epipe and VPLS services under BGP EVPN
- intra-AS BGP VPRN for VPN-IPv4 and VPN-IPv6 prefixes with both autobind and explicit SDP
- inter-AS option C for VPN-IPv4 and VPN-IPv6 VPRN prefix resolution
- multicast over IES/VPRN spoke interface with spoke SDP over an SR-TE LSP

3.12.11 Data path support

The support of SR-TE in the data path requires that the ingress LER push a label stack where each label represents a hop, a TE link, or a node, in the ERO for the LSP path computed by the router or the PCE. However, only the label and the outgoing interface to the first strict or loose hop in the ERO factor into the forwarding decision of the ingress LER. In other words, the SR-TE LSP only needs to track the reachability of the first strict or loose hop.

The first strict or loose hop of the SR-TE LSP is represented as an NHFLE to the SR shortest path tunnel. The rest of the SR-TE label stack can have a larger size and is modeled as another NHLFE referred to as a "super NHLFE".

Therefore, an SR-TE LSP is modeled in the ingress LER data path as a hierarchical LSP with the super NHLFE tunneled over the NHLFE of the SR shortest path tunnel to the first strict or loose hop in the SR-TE LSP path ERO.

Some characteristics of this design are as follows:

- The design saves on NHLFE usage. When many SR TE LSPs are going to the same first hop, they will be using the same SR shortest path tunnel and will consume one super NHLFE each, but they will be pointing to a single NHLFE, or set of NHLFEs, when ECMP exists for the first strict or loose hop, of the first-hop SR tunnel.

The ingress LER does not need to program a separate backup super NHLFE. Instead, the single super NHLFE will automatically begin forwarding packets over the LFA backup path of the SR tunnel to the first hop as soon as it is activated.

- There is an exception to the above model in the case where the user configured an empty path SR-TE LSP that uses the router's hop-to-label translation. In this case, the SR-TE LSP will use the NHLFE of the node SID of the destination router. The super NHLFE is null in this case.
- If the first segment is a node SID tunnel and multiple next hops exist, ECMP spraying is supported at the ingress LER.
- If the first-hop SR tunnel, node SID, or adjacency SID goes down, the SR module informs MPLS that the outer tunnel is down and MPLS brings the SR-TE LSP down and requests the SR to delete the SR-TE LSP in the IOM.

The data path behavior at the LSR and egress LER for an SR-TE LSP is similar to that of the shortest path tunnel because there is no tunnel state in these nodes. The forwarding of the packet is based on processing the incoming label stack consisting of a node SID and/or adjacency SID label. If the ILM is for a node SID and multiple next hops exist, ECMP spraying is supported at the LSR.

The link-protect LFA backup next hop for an adjacency SID can be programmed at the ingress LER and LSR nodes (as explained in [SR-TE LSP protection](#)).


A maximum of 12 labels, including all transport (including entropy), service, NGE, and OAM labels, can be pushed. The label stack size for the SR-TE LSP can be 1 to 11 labels, with a default value of 6.

The label stack size manipulation includes the following LER and LSR roles:

- LER role:
 - push up to 12 labels depending on the service type
 - pop up to 8 labels
- LSR role:
 - pop up to 5 labels and swap 1 label for a total of 6 labels
 - LSR hash of a packet with up to 10 labels

An example of the label stack pushed by the ingress LER and by an LSR acting as a PLR is illustrated in the following figure.

Figure 14: SR-TE LSP label stack programming



On node A, the user configures an SR-TE LSP to node D with a list of explicit strict hops mapping to the adjacency SID of links A-B, B-C, and C-D.

Ingress LER A programs a super NHLFE consisting of the label for the adjacency over link C-D and points it to the already programmed NHLFE of the SR tunnel of its local adjacency over link A-B. The latter NHLFE has the top label and also the outgoing interface to send the packet to.



Note: The SR-TE LSP does not consume a separate backup super NHLFE; it only points the single super NHLFE to the NHLFE of the SR shortest path tunnel it is riding. When the latter activates its backup NHLFE, the SR-TE LSP will automatically forward over it.

LSR Node B already programmed the primary NHLFE for the adjacency SID over link C-D and has the ILM with label 1001 point to it. In addition, node B will preprogram the link-protect LFA backup next hop for link B-C and point the same ILM to it.



Note: There is no super NHLFE at node B because it only deals with the programming of the ILM and primary and backup NHLFE of its adjacency SIDs and its local and remote node SIDs.

VPN service in node A forwards a packet to the VPN-IPv4 prefix X advertised by BGP peer D. The figure shows the resulting data path at each node for the primary path and for the FRR backup path at LSR B.

3.12.11.1 SR-TE LSP metric and MTU settings

The MPLS module assigns an SR-TE LSP the maximum LSP metric value of 16 777 215 when the local router provides the hop-to-label translation for its path. For an SR-TE LSP that uses local CSPF or PCE for path computation or has its control delegated to the PCE (**pce-control** enabled), the latter will return the computed LSP IGP or TE metric in the PCReq and PCUpd messages. In all cases, the user can override the returned value by configuring an admin metric using the command **config>router>mpls>lsp>metric**.

The MTU setting of an SR-TE LSP is derived from the MTU of the outgoing SR shortest path tunnel it is using, adjusted with the size of the super NHLFE label stack size.

The following are the details of this calculation:

$$\text{SR_Tunnel_MTU} = \text{MIN} \{ \text{Cfg_SR_MTU}, \text{IGP_Tunnel_MTU} - (1 + \text{frr-overhead}) \times 4 \}$$

where:

- **Cfg_SR_MTU** is the MTU configured by the user for all SR tunnels within a particular IGP instance using `config>router>ospf/isis>segment-routing>tunnel-mtu`. If no value was configured by the user, the SR tunnel MTU will be fully determined by the IGP interface calculation. This calculation is performed by the IGP and passed to the SR module each time it changes due to an updated resolution of the node SID.
- **IGP_Tunnel_MTU** is the minimum of the IS-IS or OSPF interface MTU among all the ECMP paths or among the primary and LFA backup paths of this SR tunnel.
- *frr-overhead* is set to 1 if **segment-routing** and **remote-lfa** options are enabled in the IGP instance; otherwise, it is set to 0.

This calculation is performed by the IGP and passed to the SR module each time it changes due to an updated resolution of the node SID.

The 7705 SAR also provides the MTU for the adjacency SID tunnel because it is needed in an SR-TE LSP if the first hop in the ERO is an adjacency SID. In that case, the calculation for **SR_Tunnel_MTU**, initially introduced for a node SID tunnel, is applied to get the MTU of the adjacency SID tunnel.

The MTU of the SR-TE LSP is derived as follows:

$$\text{SRTE_LSP_MTU} = \text{SR_Tunnel_MTU} - \text{numLabels} \times 4$$

where:

- **SR_Tunnel_MTU** is the MTU SR tunnel shortest path that the SR-TE LSP is using. The 7705 SAR also provides the MTU for the adjacency SID tunnel because it is needed in an SR-TE LSP if the first hop in the ERO is an adjacency SID. In that case, the calculation for **SR_Tunnel_MTU** (given above), initially introduced for a node SID tunnel, is applied to get the MTU of the adjacency SID tunnel.
- **numLabels** is the number of labels found in the super NHLFE of the SR-TE LSP. At LER, the super NHLFE is pointing to the SR tunnel NHLFE, which has a primary and a backup NHLFE.

This calculation is performed by the SR module and is updated each time the SR-TE LSP path changes or the SR tunnel it is using is updated.

3.12.12 SR-TE entropy labels

The 7705 SAR supports SR-TE entropy labels as described in [MPLS entropy labels](#).

3.13 Weighted ECMP by RSVP-TE or SR-TE LSPs

ECMP over MPLS LSPs implements packet spraying across multiple RSVP-TE or SR-TE LSPs within the same ECMP set.

ECMP RSVP-TE or SR-TE packet spraying consists of hashing the relevant fields in the header of a labeled packet and selecting the next-hop tunnel based on the modulo operation of the output of the hash and the number of ECMP tunnels. The maximum number of ECMP tunnels selected from the tunnel table

manager (TTM) matches the value of the user-configured **ecmp** command. Only LSPs with the same lowest LSP metric can be part of the ECMP set. If the number of these LSPs is higher than the value configured with the **ecmp** command, the LSPs with the lowest tunnel IDs are selected first.

The **ecmp** command context for setting the number of routes that are used on a service is as follows:

- for VPRN autobind weighted ECMP: **config>service>vprn>auto-bind-tunnel>ecmp max-ecmp-routes**
- for OSPF PE-CE weighted ECMP: **config>service>vprn>ecmp max-ecmp-routes**
- for VPN route-reflector forwarding with next-hop-self: **config>router>ecmp max-ecmp-routes**
- for weighted ECMP for SDPs: the number of LSPs specified in the command **config>service>sdp>lsp lspx lspy lspz** or **config>service>sdp>sr-te-lsp lspx lspy lspz**
- for EVPN autobind weighted ECMP: **config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>ecmp max-ecmp-routes** or **config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>ecmp max-ecmp-routes**



Note: The 7705 SAR supports a maximum of eight ECMP routes, which is the maximum value that an LSP can support after normalization.

With the weighted ECMP functionality, the load-balancing weight of the LSP is normalized by the system and then used to bias the amount of traffic forwarded over each LSP.

Weighted ECMP is supported on the following services:

- VPRN with autobind tunnels
- EVPN-VPWS and EVPN-VPLS with autobind tunnels
- IES and VPRN Layer 3 spoke-SDP terminated interfaces
- Epipe services with spoke SDPs
- Lpipe services with spoke SDPs
- VPLS services with spoke or mesh SDPs
- VPRN OSPF PE-CE interfaces

The weight of the LSP is configured using the **config>router>mpls>lsp>load-balancing-weight weight** and **config>router>mpls>lsp-template>load-balancing-weight weight** commands.



Note: SR-TE LSP templates are currently not supported.

3.13.1 Weighted ECMP for VPRN services with autobind

Weighted ECMP for VPRN services with autobind is enabled using the **config>service>vprn>auto-bind-tunnel>weighted-ecmp** command. This command is applicable if the autobind tunnel is configured for RSVP or SR-TE using the **config>service>vprn>auto-bind-tunnel>resolution-filter>rsvp/sr-te** command.

If weighted ECMP is enabled, a path is selected based on the output of the configured hashing algorithm. Packet paths are then mapped to LSPs for the service in proportion to the configured load-balancing weight of the LSP. The hash is based on the system load-balancing configuration. Weighted ECMP is disabled by default.

If an LSP in the ECMP set has no load-balancing weight configured, and the ECMP is set to a specific next hop, regular ECMP spraying is used.

3.13.2 Weighted ECMP for EVPN

ECMP for Layer 2 unicast traffic on Epipe and VPLS services for EVPN-MPLS destinations is enabled using the **config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>ecmp** command or **config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>ecmp** command. These commands allow the resolution of an EVPN-MPLS next hop to a group of ECMP tunnels of type RSVP-TE or SR-TE.

Weighted ECMP is enabled using the **config>service>epipe>bgp-evpn>mpls> auto-bind-tunnel>weighted-ecmp** command or **config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>weighted-ecmp** command. These commands allow packets to be sprayed across LSPs in the ECMP set according to the outcome of the hash algorithm and the configured load-balancing weight of each LSP.

EVPN VPWS ECMP and weighted ECMP behave like a non-EVPN Layer 2 Epipe and perform per-service ID hashing.

EVPN VPLS ECMP and weighted ECMP behave like a non-EVPN Layer 2 VPLS and perform per-service ID hashing but also include an internal egress ID.

3.13.3 Weighted ECMP for SDPs

Weighted ECMP for IES or VPRN Layer 3 spoke SDP interfaces, Epipe and Ipipe services with SDPs, and VPLS services with spoke or mesh SDPs is enabled using the **config>service>sdp>weighted-ecmp** command. This command is applicable when the SDP has RSVP-TE LSPs configured using the **config>service>sdp>lsp** command or SR-TE LSPs configured using the **config>service>sdp>sr-te-lsp** command.

When weighted ECMP is enabled on an SDP, a path is selected based on the configured hash. Paths are then load-balanced across the LSPs used by the SDP according to the normalized LSP load-balancing weight configured using the **load-balancing-weight** command. This means that consecutive packets of a particular service use the same LSP, but the overall load handled by LSPs to the SDP far end is balanced according to the load-balancing weight if all services using the SDP send the same bandwidth and there are more services using the SDP than there are LSPs for the SDP.

If an LSP in the ECMP set has no load-balancing weight configured, ECMP is applied to packets based on the output of the hash for the service ID.

3.14 MPLS service usage

The 7705 SAR routers enable service providers to deliver virtual private networks (VPNs) and Internet access using Generic Routing Encapsulation (GRE), IP, and/or MPLS tunnels, with Ethernet and/or SONET/SDH interfaces.

3.14.1 Service destination points


A service destination point (SDP) acts as a logical way of directing traffic from one 7705 SAR router to another through a unidirectional (one-way) service tunnel. The SDP terminates at the far-end 7705 SAR router, which directs packets to the correct service egress service access point (SAP) on that device. All services mapped to an SDP use the GRE, IP, or MPLS transport encapsulation type.

For information about service transport tunnels, see the 7705 SAR Services Guide. Service transport tunnels can support up to eight forwarding classes and can be used by multiple services.

3.15 MPLS and RSVP-TE configuration process overview

The following figure displays the process to configure MPLS and RSVP-TE parameters.

Figure 15: MPLS and RSVP-TE configuration and implementation flow



21817

3.16 Configuration notes

Network and system interfaces must be configured in the **config>router>interface** context before they can be specified in MPLS. See the 7705 SAR Router Configuration Guide for interface configuration information.

The following are MPLS and RSVP-TE guidelines and 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.
- CSPF must be enabled in order for administrative groups and SRLGs to be relevant.

3.17 Configuring MPLS and RSVP-TE with CLI

This section provides information to configure MPLS and RSVP-TE using the CLI.

Topics in this section include:

- [MPLS configuration overview](#)
- [Basic MPLS configuration](#)
- [Common configuration tasks](#)
- [MPLS configuration management tasks](#)
- [RSVP-TE configuration management tasks](#)
- [Configuring and operating SR-TE](#)

3.18 MPLS configuration overview

MPLS enables routers to forward traffic based on a label embedded in the packet header. A router examines the label to determine the next hop for the packet, instead of router address lookups to the next node when forwarding packets.

To implement MPLS on an LSP for outer tunnel and pseudowire assignment, the following entities must be configured:

- [Router interface](#)
- [Paths](#)
- [LSPs](#)
- [Pseudowires](#)
- [Signaling protocol](#) (for RSVP-TE or LDP)

3.18.1 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.

3.18.1.1 E-LSP for differentiated services

An EXP-inferred LSP (E-LSP) is an LSP that can support a variety of VLLs or traffic types. Up to eight types of traffic can be multiplexed over an E-LSP.

The prioritization of mission-critical traffic is handled by the settings of the three EXP bits. The EXP bits designate the importance of a particular packet. The classification and queuing at the provider (P) or provider edge (PE) nodes typically take place based on the value of the EXP bits. See the 7705 SAR

Quality of Service Guide for more information about the use of EXP bits and differentiated services on the 7705 SAR.

3.18.2 Paths

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

3.18.3 LSPs

The 7705 SAR supports static and dynamic LSPs.

To configure MPLS-signaled (dynamic) LSPs, the LSP must run from an ingress LER to an egress LER. Configure the dynamic LSP only at the ingress router, and either configure the LSP to allow the router software to make the forwarding decisions or configure some or all routers in the LSP path statically. The LSP is set up by RSVP-TE signaling messages. The 7705 SAR automatically manages label values. Labels that are automatically assigned have values ranging from 1024 through 1 048 575 (see [Label values](#)).

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

To establish a static LSP, an LSP must be configured from an ingress LER to an egress LER. Labels must be manually assigned and the label values must be within the range of 32 to 1023 (see [Label values](#)).

3.18.4 Pseudowires

To configure PW/VLL labels, the PW/VLL service must be configured. PW/VLL labels can be configured manually as statically allocated labels using any unused label within the static label range. Pseudowire/VLL labels can also be dynamically assigned by targeted LDP. Statically allocated labels and dynamically allocated labels are designated differently in the label information base.

PW/VLL labels are uniquely identified against a 7705 SAR, not against an interface or module.

As defined in RFC 3036, *LDP Specification*, and RFC 4447 *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*, label distribution is handled in the Downstream Unsolicited (DU) mode. Generic Label TLV is used for all setup and maintenance operations.

3.18.5 Signaling protocol

For static LSPs, the path and the label mappings and actions configured at each hop must be specified manually. RSVP-TE or LDP is not required for static LSPs.

For dynamic LSPs, RSVP-TE or LDP must be turned on. See [RSVP and RSVP-TE](#) or [Label Distribution Protocol](#).

To implement dynamic pseudowire/VLL labels, entities must be enabled as follows:

- MPLS must be enabled on all routers that are part of a static LSP
- LDP must be enabled on the ingress and egress LERs

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

3.19 Basic MPLS configuration

To enable MPLS on a 7705 SAR router, you must configure at least one MPLS interface. The MPLS interface is configured in the **config>router> mpls** context. The other MPLS configuration parameters are optional.

The following example displays an MPLS configuration output. 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.

```
A:ALU-1>config>router>if-attr# info
```

```
-----
admin-group "green" value 15
admin-group "yellow" value 20
admin-group "red" value 25
-----
```

```
A:ALU-1>config>router>mpls# info
```

```
-----
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 "to-NYC"
  hop 1 10.10.10.104 strict
  no shutdown
exit
path "secondary-path"
  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
```

3.20 Common configuration tasks

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

- [Configuring global MPLS parameters](#)
- [Configuring an MPLS interface](#)
- [Configuring MPLS paths](#)
- [Configuring an MPLS LSP](#)
- [Configuring a static LSP](#)
- [Configuring manual bypass tunnels](#)
- [Configuring RSVP-TE parameters and interfaces](#)
- [Configuring RSVP-TE message pacing parameters](#)

The following protocols must be enabled on each participating router:

- MPLS
- RSVP-TE (for RSVP-TE-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 the path 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 settings under the LSP hierarchy are optional.

3.20.1 Configuring global MPLS parameters

Admin groups can signify link colors, such as red, yellow, or green, or some other link quality. Shared risk link groups (SRLGs) are lists of interfaces that share the same risk of failure due to shared resources. MPLS interfaces advertise the admin groups and SRLGs that they support. CSPF uses the information when paths are computed for constraint-based LSPs. CSPF must be enabled in order for admin groups and SRLGs to be relevant.

Admin groups and SRLGs are created in the **config>router>if-attribute** context. Other global parameters are created in the **config>router>mpls** context.

To configure global MPLS parameters, enter the following commands:

CLI syntax:

```
config>router>if-attribute
  admin-group group-name value group-value
  srlg-group group-name value group-value
```

CLI syntax:

```
config>router>mpls
  bypass-resignal-timer minutes
  dynamic-bypass [enable | disable]
  frr-object
  hold-timer seconds
  resignal-timer minutes
  srlg-frr [strict]
```

Example:

```
config>router# if-attribute
config>router>if-attr# admin-group "green" value 15
config>router>if-attr# admin-group "red" value 25
config>router>if-attr# admin-group "yellow" value 20
config>router>if-attr# srlg-group "SRLG_fiber_1" value 50
config>router>if-attr# exit
config>router# mpls
config>router>mpls# frr-object
config>router>mpls# bypass-resignal-timer 120
config>router>mpls# hold-timer 3
config>router>mpls# resignal-timer 500
config>router>mpls# srlg-frr strict
```

The following example displays a global MPLS configuration output.

```
A:ALU-1>config>router>if-attr# info
-----
  admin-group "green" value 15
  admin-group "red" value 25
  admin-group "yellow" value 20
  srlg-group "SRLG_fiber_1" value 50
-----

A:ALU-1>config>router>mpls# info
-----
  frr-object
  bypass-resignal-timer 120
  hold-timer 3
  resignal-timer 500
  srlg-frr strict
-----

A:ALU-1>config>router>mpls# info
```

3.20.2 Configuring an MPLS interface

The interface must exist in the system before it can be configured as an MPLS interface; see the 7705 SAR Router Configuration Guide for more information.

When the MPLS protocol instance is created, the **no shutdown** command is not required because MPLS is administratively enabled upon creation. Configure the **label-map** parameters if the interface is used in a static LSP.

Use the following CLI syntax to configure an MPLS interface on a router:

CLI syntax:

```
config>router>mpls
  interface ip-int-name
    admin-group group-name [group-name...(up to 5 max)]
    label-map in-label
      pop
      swap out-label next-hop ip-address
      no shutdown
    srlg-group group-name [group-name...(up to 5 max)]
    te-metric value
    no shutdown
```

Example:

```
config>router# mpls
config>router>mpls# interface to-104
config>router>mpls>if# label-map 35
config>router>mpls>if>label-map# swap 36 next-hop 10.10.10.91
config>router>mpls>if>label-map# no shutdown
config>router>mpls>if>label-map# exit
config>router>mpls>if# srlg-group "SRLG_fiber_1"
config>router>mpls>if# no shutdown
config>router>mpls# exit
```

The following example displays the interface configuration output.

```
A:ALU-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
        srlg-group "SRLG_fiber_1"
        exit
      exit
      no shutdown
```

3.20.3 Configuring MPLS paths

When configuring an MPLS path for LSPs, the IP address of each hop 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 other routers (loose).

Use the following CLI syntax to configure a path:

CLI syntax:

```
config>router>mpls
  path path-name
    hop hop-index ip-address {strict|loose}
```

```
no shutdown
```

The following example displays a path configuration output.

```
A:ALU-1>config>router>mpls# info
-----
    interface "system"
    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
    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
-----
A:ALU-1>config>router>mpls#
```

3.20.4 Configuring an MPLS LSP

When configuring an LSP, you must specify the IP address of the egress router in the **to** statement. You must also 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.

The following displays an MPLS LSP configuration.

```
A:ALU-1>config>router>mpls>lsp# info
-----
...
    lsp "lsp-to-eastcoast"
        to 192.0.2.41
        rsvp-resv-style ff
        path-computation-method local-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:ALU-1>config>router>mpls#
```

3.20.5 Configuring a static LSP

An LSP can be explicitly (manually) configured. The reserved range of static LSP labels is 32 to 1023. Static LSPs are configured on every node along the LSP path. The label's forwarding information includes the address of the next hop router.

Use the following CLI syntax to configure a static LSP:

CLI syntax:

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

Example:

```
config>router# mpls
config>router>mpls# static-lsp static-LSP
config>router>mpls>static-lsp$ to 10.10.10.124
config>router>mpls>static-lsp# push 60 nexthop 10.10.42.3
config>router>mpls>static-lsp# no shutdown
config>router>mpls>static-lsp# exit
```

The following example displays the static LSP configuration output.

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

3.20.5.1 Configuring a fast-retry timer for static LSPs

A fast-retry timer can be configured for static LSPs. When a static LSP is trying to come up, MPLS tries to resolve the ARP entry for the next hop of the LSP. This request may fail because the next hop could still be down or unavailable. In that case, MPLS starts a retry timer before making the next request. The fast-retry command allows the user to configure the retry timer so that the LSP comes up shortly after the next hop is available.

Use the following CLI syntax to configure a fast-retry timer for static LSPs:

CLI syntax:

```
config>router>mpls
  static-lsp-fast-retry seconds
```


Example:

```
config>router# mpls
config>router>mpls# static-lsp-fast-retry 15
```

3.20.6 Configuring manual bypass tunnels

Consider the following network setup in the following figure. Assume that a manual bypass tunnel must be configured on Node B.

Figure 16: Manual bypass tunnels



20123

1. Disable dynamic bypass tunnels on Node B.
The CLI syntax for this configuration is:
config>router>mpls>dynamic-bypass [disable | enable]
By default, dynamic bypass tunnels are enabled.
2. Configure an LSP on Node B, such as B-E-F-C, which will be used only as a bypass. Specify each hop in the path and assign its **strict** or **loose** option; in this case, the bypass LSP will have a strict path. Designate the LSP as a primary LSP.
The CLI syntax for this configuration is:
config>router>mpls>path path-name >hop hop-index ip-address [strict | loose]
config>router>mpls>lsp lsp-name bypass-only
(see also the configuration example below)
Including the **bypass-only** keyword disables some options under the LSP configuration. See [Table 8: Disabled and enabled options for bypass-only](#).

Table 8: Disabled and enabled options for bypass-only

Disabled options	Enabled options
<ul style="list-style-type: none">bandwidthfast-reroutesecondary	<ul style="list-style-type: none">adaptiveadspecexcludehop-limitincludemetric-typepath-computation-method local-cspf

3. Configure an LSP from A to D and indicate fast-reroute bypass protection by selecting facility as the FRR method.

The CLI syntax for this configuration is:

config>router>mpls>lsp *lsp-name*>fast-reroute facility

If the LSP from A to D goes through Node B and bypass is requested, the next hop is Node C, and there is a manually configured bypass-only tunnel from B to C that excludes link BC (that is, path BEFC), then Node B uses the bypass-only tunnel.

The following example displays a bypass tunnel configuration output.

```
A:ALU-48>config>router>mpls># info
-----
...
    path "BEFC"
        hop 10 10.10.10.11 strict
        hop 10 10.10.10.12 strict
        hop 10 10.10.10.13 strict
        no shutdown
    exit
    lsp "bypass-BC" bypass-only
        to 10.10.10.15
        primary "BEFC"
        exit
        no shutdown
...
-----
```

3.20.7 Configuring RSVP-TE parameters and interfaces

RSVP-TE is used to set up LSPs. RSVP-TE must be enabled on the router interfaces that are participating in signaled LSPs. The default values can be modified in the **config>router>rsvp** context.

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

The following example displays an RSVP-TE configuration output.

```
A:ALU-1>config>router>rsvp# info
-----
    keep-multiplier 3
    refresh-time 30
    no msg-pacing
    rapid-retransmit-time 5
    rapid-retry-limit 3
    refresh-reduction-over-bypass disable
    no graceful-shutdown
    no entropy-label-capability
    node-id-in-rro exclude
    interface "system"
        no shutdown
    exit
    interface to-104
        hello-interval 4000
        no shutdown
    exit
    no shutdown
```

```
-----
A:ALU-1>config>router>rsvp#
```

3.20.8 Configuring RSVP-TE message pacing parameters

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

Use the following CLI syntax to configure RSVP-TE message pacing parameters:

CLI syntax:

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

The following example displays an RSVP-TE message pacing configuration output.

```
A:ALU-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:ALU-1>config>router>rsvp#
```

3.21 MPLS configuration management tasks

This section discusses the following MPLS configuration management tasks:

- [Deleting MPLS](#)
- [Modifying MPLS parameters](#)
- [Modifying an MPLS LSP](#)
- [Modifying MPLS path parameters](#)
- [Modifying MPLS static LSP parameters](#)
- [Deleting an MPLS interface](#)

3.21.1 Deleting MPLS

The **no** form of the **mpls** command typically removes an MPLS instance and all associated information. However, MPLS must be disabled (shut down) and all SDP bindings to LSPs removed before an MPLS instance can be deleted. Once MPLS is shut down, the **no mpls** command deletes the protocol instance and removes all configuration parameters for the MPLS instance.

If MPLS is not shut down first, when the **no mpls** command is executed, a warning message on the console indicates that MPLS is still administratively up.

To delete the MPLS instance:

1. Disable the MPLS instance using the **shutdown** command.
2. Remove the MPLS instance from the router using the **no mpls** command.

CLI syntax:

```
config>router# no mpls
```

3.21.2 Modifying MPLS parameters



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

3.21.3 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.

The following example displays an MPLS LSP configuration output. See [Configuring an MPLS interface](#).

```
A:ALU-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:ALU-1>config>router>mpls#
```

3.21.4 Modifying MPLS path parameters

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

The following example displays an MPLS path configuration output. See [Configuring MPLS paths](#).

```
A:ALU-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
#-----
```

3.21.5 Modifying MPLS static LSP parameters

Use the **show>service>router>static-lsp** command to display a list of LSPs.

To modify static LSP parameters, the **config>router>mpls>static-lsp *lsp-name*** context must be shut down.

To modify an LSP:

1. Access the specific LSP by specifying the LSP name, and then shut it down.
2. Enter the parameter to modify and then enter the new information.

Example:

```
config>router# mpls
config>router>mpls# static-lsp "static-LSP"
config>router>mpls>static-lsp# shutdown
config>router>mpls>static-lsp# to 10.10.0.234
config>router>mpls>static-lsp# push 1023 nexthop 10.10.8.114
config>router>mpls>static-lsp# no shutdown
config>router>mpls>static-lsp# exit
```

The following example displays the static LSP configuration output.

```
ALU-1>config>router>mpls# info
#-----
...
    static-lsp "static-LSP"
        to 10.10.10.234
        push 1023 nexthop 10.10.8.114
        no shutdown
    exit
    no shutdown
#-----
ALU-1>config>router>mpls#
```

3.21.6 Deleting an MPLS interface

To delete an interface from the MPLS configuration:

1. Administratively disable the interface using the **shutdown** command.
2. Delete the interface with the **no interface** command.

CLI syntax:

```
mpls
  interface ip-int-name
    shutdown
  exit
no interface ip-int-name
```

Example:

```
config>router# mpls
config>router>mpls# interface to-104
config>router>mpls>if# shutdown
config>router>mpls>if# exit
config>router>mpls# no interface to-104
```

The following example displays the configuration output when interface "to-104" has been deleted.

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

3.22 RSVP-TE configuration management tasks

This section discusses the following RSVP-TE configuration management tasks:

- [Modifying RSVP-TE parameters](#)
- [Modifying RSVP-TE message pacing parameters](#)
- [Deleting an interface from RSVP-TE](#)

3.22.1 Modifying RSVP-TE parameters

Only interfaces configured in the MPLS context can be modified in the **rsvp** context.

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

The following example displays a modified RSVP-TE configuration output.

```
A:ALU-1>config>router>rsvp# info
-----
    keep-multiplier 5
    refresh-time 60
    msg-pacing
        period 400
        max-burst 400
    exit
    rapid-retransmit-time 5
    rapid-retry-limit 3
    refresh-reduction-over-bypass disable
    no graceful-shutdown
    no entropy-label-capability
    no implicit-null-label
    node-id-in-rro exclude
    interface "system"
    exit
    interface "test1"
        hello-interval 5000
    exit
    no shutdown
-----
A:ALU-1>config>router>rsvp#
```

3.22.2 Modifying RSVP-TE message pacing parameters

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

The following example displays a modified RSVP-TE message pacing configuration output. See [Configuring RSVP-TE message pacing parameters](#).

```
A:ALU-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:ALU-1>config>router>rsvp#
```

3.22.3 Deleting an interface from RSVP-TE

Interfaces cannot be deleted directly from the RSVP-TE configuration. Because an interface is created in the **mpls** context and then configured in the **rsvp** context, it can only be deleted in the mpls context. This removes the association from RSVP-TE.

See [Deleting an MPLS interface](#).

3.23 Configuring and operating SR-TE

This section provides information about the configuration and operation of the Segment Routing with Traffic Engineering (SR-TE) LSP. It covers the following topics:

- [SR-TE configuration prerequisites](#)
- [SR-TE LSP configuration overview](#)
- [Configuring path computation and control for SR-TE LSPs](#)
- [Configuring SR-TE LSP label stack size](#)
- [Configuring adjacency SID parameters](#)
- [Configuring PCC-controlled, PCE-computed, and PCE-controlled SR-TE LSPs](#)

3.23.1 SR-TE configuration prerequisites

To configure SR-TE, the user must first configure the prerequisite parameters.

First, configure the label space partition for the segment routing global block (SRGB) for all participating routers in the segment routing domain by using the **mpls-labels>sr-labels** command.

Example:

```
configure>router>mpls-labels
  sr-labels start 200000 end 200400
exit
```

Enable segment routing, traffic engineering, and advertisement of router capability in all participating IGP instances in all participating routers by using the **traffic-engineering**, **advertise-router-capability**, and **segment-routing** commands.

Example:

```
ospf 0
  traffic-engineering
  advertise-router-capability area
  loopfree-alternates remote-lfa
  area 0.0.0.202
    stub
      no summaries
    exit
  interface "system"
    node-sid index 194
    no shutdown
  exit
  interface "toSim199"
    interface-type point-to-point
    no shutdown
  exit
  interface "toSim213"
    interface-type point-to-point
    no shutdown
  exit
  interface "toSim219"
    interface-type point-to-point
    metric 2000
    no shutdown
  exit
```

```

exit
segment-routing
    prefix-sid-range global
    entropy-label {force-disable | enable}
    no shutdown
exit
no shutdown
exit

```

Configure a segment routing tunnel MTU for the IGP instance, if required, by using the **tunnel-mtu** command.

Example:

```

prefix-sid-range global
tunnel-mtu 1500
no shutdown

```

Assign a node SID to each loopback interface that a router would use as the destination of a segment routing tunnel by using the **node-sid** command.

Example:

```

ospf 0
    area 0.0.0.202
        interface "system"
            node-sid index 194
            no shutdown
        exit

```

3.23.2 SR-TE LSP configuration overview

An SR-TE LSP can be configured as an LSP using the existing CLI command hierarchy under the MPLS context and specifying the **sr-te** LSP type.

CLI syntax:

```

config>router>mpls>lsp lsp-name | sr-te

```

A primary path can be configured for the SR-TE LSP.

Use the following CLI syntax to associate an empty path or a path with strict or loose explicit hops with the primary paths of the SR-TE LSP:

CLI syntax:

```

config>router>mpls>path>hop hop-index ip-address {strict | loose}
config>router>mpls>lsp>primary path-name

```

3.23.3 Configuring path computation and control for SR-TE LSPs

Use the following syntax to configure the path computation requests only (PCE-computed) or both path computation requests and path updates (PCE-controlled) to the PCE for a specific LSP:

CLI syntax:

```

config>router>mpls>lsp>path-computation-method pce

```

```
config>router>mpls>lsp>pce-control
```

The PCC LSP database is synchronized with the PCE LSP database using the PCEP PCRpt (PCE Report) message for LSPs that have the following commands enabled:

CLI syntax:

```
config>router>mpls>pce-report sr-te {enable | disable}  
config>router>mpls>lsp>pce-report {enable | disable | inherit}
```

3.23.3.1 Configuring path profile and group for PCC-initiated and PCE-computed/controlled LSPs

The PCE supports the computation of disjoint paths for two different LSPs originating or terminating on the same or different PE routers. To indicate this constraint to the PCE, the user must configure the PCE path profile ID and path group ID that the LSP belongs to. These parameters are passed transparently by the PCC to the PCE and are therefore opaque data to the router. Use the following syntax to configure the path profile and path group:

CLI syntax:

```
config>router>mpls>lsp>path-profile profile-id [path-group group-id]
```

The association of the optional path group ID is to allow the PCE to determine which profile ID this path group ID must be used with. One path group ID is allowed per profile ID. The user can, however, enter the same path group ID with multiple profile IDs by executing this command multiple times. A maximum of five entries of **path-profile** [*path-group*] can be associated with the same LSP. More details of the operation of the PCE path profile are provided in the [PCEP](#) chapter.

3.23.4 Configuring SR-TE LSP label stack size

Use the following syntax to configure the maximum number of labels that the ingress LER can push for a specific SR-TE LSP:

CLI syntax:

```
config>router>mpls>lsp>max-sr-labels label-stack-size
```

This command allows the user to reduce the SR-TE LSP label stack size by accounting for additional transport (including entropy), service, and other labels when packets are forwarded in a particular context. See [Data path support](#) for more information about label stack size requirements in various forwarding contexts. If the CSPF on the PCE or the router's hop-to-label translation cannot find a path that meets the maximum SR label stack, the SR-TE LSP will remain on its current path or will remain down if it has no path. The range is 1 to 11 labels with a default value of 6.

3.23.5 Configuring adjacency SID parameters

Configure the adjacency hold timer for the LFA or remote LFA backup next hop of an adjacency SID.

Use the following syntax to configure the length time that LTN or ILM records of an adjacency SID are kept:

CLI syntax:

```
config>router>ospf>segment-routing>adj-sid-hold seconds[1..300, default 15]
config>router>isis>segment-routing>adj-sid-hold seconds[1..300, default 15]
```

Example:

```
adj-sid-hold 15
prefix-sid-range global
no tunnel-table-pref
no tunnel-mtu
no backup-node-sid
no shutdown
```

While protection is enabled globally for all node SIDs and local adjacency SIDs when the user enables the **loopfree-alternates** option in IS-IS or OSPF at the LER and LSR, there are applications where the user wants traffic to never divert from the strict hop computed by CSPF for an SR-TE LSP. In that case, use the following syntax to disable protection for all adjacency SIDs formed over a particular network IP interface:

CLI syntax:

```
config>router>ospf>area>if>no sid-protection
config>router>isis>if>no sid-protection
```

Example:

```
node-sid index 194
no sid-protection
no shutdown
```

3.23.6 Configuring PCC-controlled, PCE-computed, and PCE-controlled SR-TE LSPs

The following example shows the configuration of PCEP PCC parameters on LER routers that require peering with the PCE server:

Example:

```
keepalive 30
dead-timer 120
no local-address
unknown-message-rate 10
report-path-constraints
peer 192.0.0.226
  no shutdown
exit
no shutdown
```

The following example shows the configuration of a PCC-controlled SR-TE LSP that is not reported to the PCE:

Example:

```
lsp "to-SanFrancisco" sr-te
  to 192.0.2.211
  path-computation-method local-cspf
  pce-report disable
  metric 10
  primary "loose-anycast"
```

```
exit
no shutdown
exit
```

The following example shows the configuration of a PCC-controlled SR-TE LSP that is reported to the PCE:

Example:

```
lsp "to-SanFrancisco" sr-te
to 192.0.2.211
path-computation-method local-cspf
pce-report enable
metric 10
primary "loose-anycast"
exit
no shutdown
exit
```

The following example shows the configuration of a PCE-computed SR-TE LSP that is reported to the PCE:

Example:

```
lsp "to-SanFrancisco" sr-te
to 192.0.2.211
path-computation-method pce
pce-report enable
metric 10
primary "loose-anycast"
exit
no shutdown
exit
```

The following example shows the configuration of a PCE-controlled SR-TE LSP with no PCE path profile:

Example:

```
lsp "from Reno to Atlanta no Profile" sr-te
to 192.0.2.224
path-computation-method pce
pce-report enable
pce-control
primary "empty"
exit
no shutdown
exit
```

The following example shows the configuration of a PCE-controlled SR-TE LSP with a PCE path profile and a maximum label stack set to a non-default value:

Example:

```
lsp "from Reno to Atlanta no Profile" sr-te
to 192.0.2.224
max-sr-labels 8 additional-frr-labels 1
path-computation-method pce
pce-report enable
pce-control
path-profile 10 path-group 2
primary "empty"
bandwidth 15
exit
no shutdown
```



```
exit
```

3.24 MPLS and RSVP-TE command reference

3.24.1 Command hierarchies

- [MPLS commands](#)
- [RSVP-TE commands](#)
- [Show commands](#)
- Tools commands (see the Tools section of the 7705 SAR OAM and Diagnostics Guide)
- [Clear commands](#)
- [Monitor commands](#)
- [Debug commands](#)

3.24.1.1 MPLS commands

```

config
- router [router-name]
- [no] mpls
- [no] admin-group-frr
- auto-lsp lsp-template template-name {policy peer-prefix-policy [peer-prefix-
policy...(up to 5 max)] | one-hop}
- no auto-lsp lsp-template template-name
- bypass-resignal-timer minutes
- no bypass-resignal-timer
- [no] cspf-on-loose-hop
- dynamic-bypass [enable | disable]
- entropy-label rsvp-te {force-disable | enable}
- entropy-label sr-te {force-disable | enable}
- [no] frr-object
- hold-timer seconds
- no hold-timer
- ingress-statistics
- [no] lsp lsp-name sender ip-address
- accounting-policy policy-id
- no accounting-policy
- [no] collect-stats
- [no] shutdown
- [no] interface ip-int-name
- [no] admin-group group-name [group-name...(up to 5 max)]
- [no] label-map in-label
- [no] pop
- swap out-label nexthop ip-address
- no swap
- [no] shutdown
- [no] shutdown
- [no] srlg-group group-name [group-name...(up to 5 max)]
- te-metric value
- no te-metric
- least-fill-min-thd percent
- no least-fill-min-thd
- least-fill-reoptim-thd percent
- no least-fill-reoptim-thd
- [no] logger-event-bundling

```

```

- [no] lsp lsp-name [bypass-only] [sr-te]
- [no] adaptive
- [no] adspec
- bfd
  - [no] bfd-enable
  - bfd-template name
  - no bfd-template
  - failure-action failover-or-down
  - no failure-action
  - return-path-label label-value
  - no return-path-label
  - wait-for-up-timer seconds
  - no wait-for-up-timer
- bgp-transport-tunnel {include | exclude}
- [no] egress-statistics
  - accounting-policy policy-id
  - no accounting-policy
  - [no] collect-stats
  - [no] shutdown
- entropy-label {force-disable | inherit | enable}
- [no] exclude group-name [group-name...(up to 5 max)]
- fallback-path-computation-method [local-cspf | none]
- no fallback-path-computation-method
- [no] fast-reroute [frr-method]
  - hop-limit limit
  - no hop-limit
  - [no] node-protect
  - [no] propagate-admin-group
- from ip-address
- hop-limit number
- no hop-limit
- igp-shortcut [lfa-protect | lfa-only | relative-metric [offset]]
- no igp-shortcut
- [no] include group-name [group-name...(up to 5 max)]
- [no] label-stack-reduction
- [no] least-fill
- load-balancing-weight weight
- no load-balancing-weight
- local-sr-protection local-sr-protection
- no local-sr-protection
- max-sr-labels label-stack-size [additional-frr-labels labels]
- no max-sr-labels
- metric metric
- metric-type metric-type
- no metric-type
- path-computation-method path-computation-method
- no path-computation-method
- path-profile profile-id [path-group group-id]
- no path-profile profile-id
- pce-associations
  - [no] diversity diversity-assoc-name
  - [no] policy policy-assoc-name
- [no] pce-control
- pce-report {enable | disable | inherit}
- [no] primary path-name
  - [no] adaptive
  - bandwidth rate-in-mps
  - no bandwidth
  - bfd
    - [no] bfd-enable
    - bfd-template name
    - no bfd-template
    - return-path-label label-value
    - no return-path-label

```

```

- wait-for-up-timer seconds
- no wait-for-up-timer
- [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] priority setup-priority
- [no] record
- [no] record-label
- [no] shutdown
- [no] propagate-admin-group
- 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-mbps
- no bandwidth
- bfd
- [no] bfd-enable
- bfd-template name
- no bfd-template
- return-path-label label-value
- no return-path-label
- wait-for-up-timer seconds
- no wait-for-up-timer
- [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)]
- path-preference preference-number
- no path-preference
- [no] priority setup-priority
- [no] record
- [no] record-label
- [no] shutdown
- [no] srlg
- [no] standby
- [no] shutdown
- to ip-address
- vprn-auto-bind [include | exclude]
- no vprn-auto-bind
- lsp-template template-name
- lsp-template template-name mesh-p2p
- lsp-template template-name one-hop-p2p
- no lsp-template template-name [one-hop-p2p | mesh-p2p]
- [no] adaptive
- [no] adspec
- bgp-transport-tunnel {include | exclude}
- cspf [use-te-metric]
- no cspf
- [no] default-path path-name
- [no] exclude group-name [group-name...(up to 5 max)]
- [no] fast-reroute [frr-method]
- hop-limit limit
- no hop-limit
- [no] node-protect
- [no] propagate-admin-group
- from ip-address
- hop-limit number
- no hop-limit
- igp-shortcut [lfa-protect | lfa-only] [relative-metric [offset]]

```

```

- no igp-shortcut
- [no] include group-name [group-name...(up to 5 max)]
- [no] least-fill
- load-balancing-weight weight
- no load-balancing-weight
- metric metric
- pce-report {enable | disable | inherit}
- [no] propagate-admin-group
- retry-limit number
- no retry-limit
- retry-timer seconds
- no retry-timer
- [no] shutdown
- vprn-auto-bind [include | exclude]
- [no] path path-name
- hop hop-index ip-address {strict | loose}
- hop hop-index sid-label sid-value
- no hop hop-index
- [no] shutdown
- pce-report rsvp-te {enable | disable}
- pce-report sr-te {enable | disable}
- resignal-timer minutes
- no resignal-timer
- srlg-frr [strict]
- no srlg-frr
- sr-te-resignal-timer minutes
- no sr-te-resignal-timer
- [no] shutdown
- [no] static-lsp lsp-name
- push label nexthop ip-address
- no push label
- to ip-address
- [no] shutdown
- static-lsp-fast-retry seconds
- no static-lsp-fast-retry
- mpls-labels
- [no] reserved-label-block name
- start-label start-value end-label end-value
- no start-label
- sr-labels start start-value end end-value
- no sr-labels
- static-label-range static-range
- no static-label-range

```

3.24.1.2 RSVP-TE commands

```

config
- router
- [no] rsvp
- [no] entropy-label-capability
- [no] graceful-shutdown
- [no] implicit-null-label
- [no] interface ip-int-name
- auth-keychain name
- no auth-keychain
- authentication-key {authentication-key | hash-key} [hash | hash2]
- no authentication-key
- [no] bfd-enable
- [no] graceful-shutdown
- hello-interval milli-seconds
- no hello-interval

```

```

- implicit-null-label {enable | disable}
- no implicit-null-label
- [no] refresh-reduction
  - [no] reliable-delivery
- [no] shutdown
- subscription percentage
- no subscription
- [no] keep-multiplier number
- no keep-multiplier
- [no] msg-pacing
  - max-burst number
  - no max-burst
  - period milli-seconds
  - no period
- node-id-in-rro {include | exclude}
- rapid-retransmit-time hundred-milliseconds
- no rapid-retransmit-time
- rapid-retry-limit number
- no rapid-retry-limit
- refresh-reduction-over-bypass {enable | disable}
- refresh-time seconds
- no refresh-time
- [no] shutdown

```

3.24.1.3 Show commands

```

show
- router
  - mpls
    - admin-group group-name
    - bypass-tunnel [to ip-address] [protected-lsp [lsp-name]] [dynamic | manual]
  [detail]
    - interface [ip-int-name | ip-address] [label-map [label]]
    - interface [ip-int-name | ip-address] statistics
    - lsp [lsp-name] [status {up | down}] [from ip-address | to ip-address] [detail]
  [auto-lsp {all | mesh-p2p | one-hop-p2p}]
    - lsp {transit | terminate} [status {up | down}] [from ip-address | to ip-address |
lsp-name name] [detail]
      - lsp count
      - lsp [lsp-name] activepath [auto-lsp {all | mesh-p2p | one-hop-p2p}]
      - lsp [lsp-name] path [path-name] [status {up | down}] [detail] [auto-lsp {all |
mesh-p2p | one-hop-p2p}]
      - lsp [lsp-name] path [path-name] mbb [auto-lsp {all | mesh-p2p | one-hop-p2p}]
      - lsp-egress-stats [type lsp-type] [active] [template-match]
      - lsp-egress-stats lsp lsp-name
      - lsp-ingress-stats [type lsp-type] [active] [template-match SessionNameString
[sender ip-address]]
      - lsp-ingress-stats lsp lsp-name sender ip-address
      - lsp-template [lsp-template-name] bindings
      - lsp-template [lsp-template-name] detail
      - path [path-name] [lsp-binding]
      - sr-te-lsp [lsp-name] [status {up | down}] [detail] path [path-name]
      - sr-te-lsp [lsp-name] [detail]
      - sr-te-lsp [sp-name] [status {up | down}] [to ip-address] [detail]
      - static-lsp [lsp-name]
      - static-lsp [lsp-type]
      - static-lsp count
      - status
show
- router
  - mpls-labels

```

```

- label start-label [end-label | in-use | label-owner]
- label-range
- summary
show
- router
  - bfd
    - seamless-bfd
    - session lsp-index lsp-index [path-lspid path-lspid] [detail]
    - session lsp-name lsp-name [path-lspid path-lspid] [detail]
    - session lsp-path [detail]
    - session lsp-path prefix ip-prefix/prefix-length [src ip-address] [detail]
  - rsvp
    - interface [ip-int-name | ip-address] statistics [detail]
    - neighbor [ip-address] [detail]
    - session [session-type] [from ip-address | to ip-address | lsp-name name] [status
{up | down}] [detail]
    - statistics
    - status

```

3.24.1.4 Clear commands

```

clear
- router
  - mpls
    - interface [ip-int-name] [statistics]
    - lsp [lsp-name]
    - lsp-egress-stats [lsp-name]
    - lsp-ingress-stats [ip-address lsp lsp-name]
  - rsvp
    - interface [ip-int-name] [statistics]
    - statistics

```

3.24.1.5 Monitor commands

```

monitor
- router
  - mpls
    - interface interface [interface...[(up to 5 max)] [interval seconds]
[repeat repeat] [absolute | rate]
    - lsp-egress-stats lsp lsp-name [interval seconds] [repeat repeat] [absolute |
rate]
    - lsp-ingress-stats lsp lsp-name [interval seconds] [repeat repeat] [absolute |
rate] ip-address
  - rsvp
    - interface interface [interface...[(up to 5 max)] [interval seconds]
[repeat repeat] [absolute | rate]

```

3.24.1.6 Debug commands

```

debug
- router
  - [no] mpls [lsp lsp-name] [sender source-address] [endpoint endpoint-address] [tunnel-
id tunnel-id] [lsp-id lsp-id] [interface ip-int-name]
  - [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
- [no] rsvp [lsp lsp-name] [sender sender-address] [endpoint endpoint-address] [tunnel-
id tunnel-id] [lsp-id lsp-id] [interface ip-int-name]
- [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
- ack [detail]
- no ack
- all [detail]
- no all
- bundle [detail]
- no bundle
- hello [detail]
- no hello
- path [detail]
- no path
- patherr [detail]
- no patherr
- pathtear [detail]
- no pathtear
- resv [detail]
- no resv
- resvterr [detail]
- no resvterr
- resvtear [detail]
- no resvtear
- srefresh [detail]
- no srefresh

```


3.24.2 Command descriptions

- [Configuration commands \(MPLS\)](#)
- [Configuration commands \(RSVP-TE\)](#)
- [Show commands \(MPLS\)](#)
- [Show commands \(MPLS-labels\)](#)
- [Show commands \(seamless BFD\)](#)
- [Show commands \(RSVP\)](#)
- [Clear commands](#)
- [Monitor commands](#)
- [Debug commands](#)

3.24.2.1 Configuration commands (MPLS)

- [Generic commands](#)
- [MPLS-labels commands](#)
- [MPLS global commands](#)
- [RSVP LSP statistics commands](#)
- [Interface commands](#)
- [Interface label-map commands](#)
- [LSP and LSP template commands](#)
- [Primary and secondary path commands](#)
- [LSP path commands](#)
- [Static LSP commands](#)

3.24.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

config>router>mpls

config>router>mpls>ingress-statistics>lsp

config>router>mpls>interface

config>router>mpls>if>label-map

```
config>router>mpls>lsp>egress-statistics
config>router>mpls>path
config>router>mpls>static-lsp
```

Description

The **shutdown** command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many objects must be shut down before they can be deleted. Many entities must be explicitly enabled using the **no shutdown** command.

In the **label-map** context, all packets that match the specified *in-label* are dropped when the label map is shut down.

In the **path** context, 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 places the entity into an administratively enabled state. In the **mpls** and **mpls>interface** contexts, this triggers any LSPs that were previously defined under the associated context to come back up. In the **path** context, the **no** form of this command administratively enables the path and all LSPs—where the path is defined as a primary or a standby secondary path—are (re)established.

Default

mpls – no shutdown
interface – shutdown
label-map – no shutdown
path – shutdown
static-lsp – shutdown

3.24.2.1.2 MPLS global commands

mpls

Syntax

[no] mpls

Context

config>router

Description

This command creates the MPLS protocol instance and enables MPLS configuration. The MPLS protocol instance is not created by default, but once it is created, a **no shutdown** command is not required since MPLS is enabled automatically. The **shutdown** command administratively disables MPLS.

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

MPLS must be shut down and all SDP bindings to LSPs removed before the MPLS instance can be deleted. If MPLS is not shut down, when the **no mpls** command is executed, a warning message on the console indicates that MPLS is still administratively up.

admin-group-frr

Syntax

[no] admin-group-frr

Context

config>router>mpls

Description

This command enables the use of admin-group constraints in the association of a manual or dynamic bypass LSP 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 in the PATH message of the LSP primary path. If the FAST_REROUTE object is not included in the PATH message, 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, it only uses the admin-group constraint from the LSP and/or path level configurations.

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

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

Changes to this command (enabling or disabling) will apply only to new attempts to find a valid bypass.

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

Default

no admin-group-frr

auto-lsp

Syntax

auto-lsp lsp-template *template-name* {policy *peer-prefix-policy* [*peer-prefix-policy*...(up to 5 max)] | one-hop}

no auto-lsp lsp-template *template-name*

Context

config>router>mpls

Description

This command enables the automatic creation of an RSVP-TE point-to-point LSP within a single IGP IS-IS level or OSPF area that can subsequently be used by services and/or IGP shortcuts. It can be used to create an RSVP-TE LSP mesh to a destination node whose router ID matches a prefix in a specified previously created peer prefix policy, or to create single-hop RSVP-TE LSPs. These LSP types are referred to as auto-LSP of type mesh or auto-LSP of type one-hop.

Multiple templates can be associated with the same or different peer prefix policies. Each application of an LSP template with a given prefix in the prefix list results in the instantiation of a single CSPF-computed LSP primary path using the LSP template parameters, as long as the prefix corresponds to a router ID for a node in the TE database. Auto-LSP does not support the automatic signaling of a secondary path for an LSP. If the signaling of multiple LSPs to the same destination node is required, a separate LSP template must be associated with a prefix list that contains the same destination node address. Each instantiated LSP will have a unique LSP ID and a unique tunnel ID. Auto-LSP also does not support the signaling of a non-CSPF LSP. The selection of the **no cspf** option in the LSP template is blocked.

Up to five peer prefix policies can be associated with an LSP template. Every time the user executes the above command with the same or different prefix policy associations or a prefix policy associated with the LSP template, the system re-evaluates the prefix policy. The outcome of the re-evaluation indicates to MPLS whether an existing LSP must be torn down or a new LSP must be signaled to a destination address that is already in the TE database.

If a /32 prefix is added to or removed from a prefix list associated with the template, or if a prefix range is expanded or narrowed, the prefix policy re-evaluation described above is performed.

A **no shutdown** of the template must be performed before it takes effect. When a template is in use, it must be shut down before the user can make any changes to the parameters except for LSP parameters for which the change can be handled with the make-before-break (MBB) procedures. This includes **fast-reroute** with or without the **hop-limit** or **node-protect** options. When the template is shut down and parameters are added, removed or modified, the existing instances of the LSP using this template are torn down and resigaled.

The trigger to signal the LSP is when the router with a router ID matching a prefix in the prefix list appears in the TE database. The signaled LSP is installed in the tunnel table manager (TTM) and is available to applications such as resolution of BGP label routes, and resolution of BGP, IGP, and static routes. It can also be used for autobinding by a VPRN service but cannot be used as a provisioned SDP for explicit binding.

Except for the MBB limitations to the configuration parameter change in the LSP template, MBB procedures for manual and timer-based resigaling of the LSP, and for TE graceful shutdown, are supported.

The **one-to-one** option under fast-reroute is not supported.

If the **one-hop** option is specified instead of a prefix policy, this command enables the automatic signaling of single-hop, point-to-point LSPs using the specified template to all directly connected neighbors. This LSP type is referred to as auto-LSP of type **one-hop**. When the above command is executed, the TE database keeps track of each TE link to a directly connected IGP neighbor whose router ID is discovered. It then instructs MPLS to signal an LSP with a destination address matching the router ID of the neighbor and with a strict hop consisting of the address of the interface used by the TE link. This results in one or more LSPs signaled to the neighboring router.

For an LSP mesh, the **no** form of this command deletes all LSPs signaled using the specified template and prefix policy. When the **one-hop** option is used, the **no** form of the command deletes all single-hop LSPs signaled using the specified template to all directly connected neighbors.

Default

n/a

Parameters

template-name

specifies an LSP template name

one-hop

specifies that the template type is one-hop LSP, rather than LSP mesh

peer-prefix-policy

specifies a peer prefix policy name. The prefix policy must already be defined.

bypass-resignal-timer**Syntax**

bypass-resignal-timer *minutes*

no bypass-resignal-timer

Context

config>router>mpls

Description

This command triggers the periodic global reoptimization of all dynamic bypass LSP paths associated with RSVP point-to-point LSPs. The operation is performed at each expiry of the timer.

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

Default

no bypass-resignal-timer

Parameters

minutes

the time 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
```

Description

This command enables the option to perform CSPF calculations to the next loose hop or the final destination of the LSP on the LSR. On receiving a PATH message on the LSR and processing all local hops in the received ERO, if the next hop is loose, then the LSR does a CSPF calculation to the next loose hop (this is known as ERO expansion). On successful completion of the CSPF calculation, the ERO in the PATH message is modified to include the 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.

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

Default

```
no cspf-on-loose-hop
```

dynamic-bypass

Syntax

```
dynamic-bypass [enable | disable]
```

Context

```
config>router>mpls
```

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.

Default

```
enable
```

entropy-label

Syntax

```
entropy-label rsvp-te {force-disable | enable}
```

```
entropy-label sr-te {force-disable | enable}
```

Context

```
config>router>mpls
```

Description

This command enables or disables the use of entropy labels for MPLS RSVP-TE and SR-TE LSPs.

If **entropy-label** is enabled, the entropy label and entropy label indicator (ELI) are inserted in the label stack. In some cases, this may result in an unsupported label stack depth or large changes in the label stack depth during the lifetime of an LSP (for example, due to switching from a primary path with entropy label capability (ELC) enabled to a secondary path for which the far end has not signaled ELC).

This command provides local control at the head end of an RSVP-TE or SR-TE LSP over whether an entropy label is inserted on the LSP by overriding the ELC signaled from the far-end LER, and control over how the additional label stack depth is accounted for.

By default, the value of **entropy-label** is inherited from the MPLS level. This command overrides the default MPLS behavior on a per-LSP basis. For auto-LSPs, it can only be configured in LSP templates of type **one-hop-p2p** and **mesh-p2p**.

When the value of **entropy-label** changes at either the MPLS level or the LSP level, the new operational value does not take effect until the LSP is resigaled. A **shutdown/no shutdown** command must be performed on the LSP to enable the new value.

The user can use the **clear** command or bounce MPLS using the **shutdown/no shutdown** command to force the new value to take effect for a large numbers of LSPs.

Default

entropy-label disable

Parameters

rsvp-te

indicates that the **entropy-label** command applies to RSVP-TE LSPs

sr-te

indicates that the **entropy-label** command applies to SR-TE LSPs

force-disable

the ingress LER will not consider the entropy label or the ELI in the label stack while sending the information to the TTM and NHLFE. The system will mark the TTM and NHLFE as ELC not supported, and applications will not insert an entropy label or ELI in the label stack.

enable

the ingress LER will take into consideration what is signaled from the egress node for ELC for marking the NHLFE, while the TTM is always marked. Although applications will only insert the entropy label if the far end signals ELC, the additional two labels of the entropy label and ELI are always accounted for.

frr-object

Syntax

[no] frr-object

Context

config>router>mpls

Description

This command specifies whether signaling the FAST_REROUTE object is on or off. 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 – by default, the value is inherited by all LSPs

hold-timer

Syntax

hold-timer *seconds*

no hold-timer

Context

config>router>mpls

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 status is up.

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

Parameters

seconds

specifies the hold time, in seconds

Values 0 to 10

least-fill-min-thd

Syntax

least-fill-min-thd *percent*

no least-fill-min-thd

Context

config>router>mpls

Description

This parameter is used in the least-fill path selection process. See the description of the [least-fill](#) command for information about the least-fill path selection process. When comparing the percentages of least

available link bandwidth across the available paths, whenever two percentages differ by less than the value configured as the least-fill minimum threshold, CSPF considers them to be equal and applies a random number generator to select the path.

The **no** form of the command resets this parameter to its default value.

Default

5

Parameters

percent

specifies the least fill minimum threshold value as a percentage

Values 1 to 100

least-fill-reoptim-thd

Syntax

least-fill-reoptim-thd *percent*

no least-fill-reoptim-thd

Context

config>router>mpls

Description

This parameter is used in the least-fill path selection process. See the description of the [least-fill](#) command for information about the least-fill path selection process.

During a timer-based resignaling of an LSP path that has the least-fill option enabled, CSPF first updates the least-available bandwidth value for the current path of this LSP. It then applies the least-fill path selection method to select a new path for this LSP. If the new computed path has the same cost as the current path, CSPF compares the least-available bandwidth values of the two paths and if the difference exceeds the user-configured optimization threshold, MPLS generates a trap to indicate that a better least-fill path is available for this LSP. This trap can be used by an external SNMP-based device to trigger a manual resignaling of the LSP path, since the timer-based resignaling will not resignal the path in this case. MPLS generates a path update trap at the first MBB event that results in the resignaling of the LSP path. This clears the eligibility status of the path at the SNMP device.

The **no** form of the command resets this parameter to its default value.

Default

10

Parameters

percent

specifies the least fill reoptimization threshold value as a percentage

Values 1 to 100

logger-event-bundling

Syntax

[no] logger-event-bundling

Context

config>router>mpls

Description

This command merges two of the most commonly generated MPLS traps, vRtrMplsXCCreate and vRtrMplsXCDelete, which can be generated at both the LER and LSR, into the new vRtrMplsSessionsModified trap. In addition, this command bundles traps of multiple RSVP sessions, such as LSPs, into this new trap.

This trap bundling allows the user to minimize trap generation in an MPLS network. MPLS trap throttling is not applied to the vRtrMplsSessionsModified trap.

The **no** version of the command disables the merging and bundling of the vRtrMplsXCCreate and vRtrMplsXCDelete traps.

pce-report

Syntax

pce-report rsvp-te {enable | disable}

pce-report sr-te {enable | disable}

Context

config>router>mpls

Description

This command separately configures the reporting modes to a PCE for RSVP-TE or SR-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.

The global MPLS-level **pce-report** command can be used to enable or disable PCE reporting for all SR-TE LSPs or RSVP-TE LSPs during PCE LSP database synchronization. This configuration is inherited by all LSPs of a particular type (RSVP-TE LSPs or SR-TE LSPs). The PCC reports both CSPF and non-CSPF LSPs.

The LSP-level **pce-report** command overrides the global configuration for the reporting of an LSP to the PCE (see **config>router>mpls>lsp>pce-report**). The default configuration is to inherit the global MPLS-level configuration.

The default configuration is disabled. This default configuration is meant to control the introduction of a PCE into an existing network and let the operator decide whether all RSVP-TE LSPs or SR-TE LSPs

need to be reported. If PCE reporting is disabled for an LSP, either due to inheritance of the global MPLS configuration or due to LSP-level configuration, enabling the **pce-control** option for the LSP has no effect.

Default

pce-report rsvp-te disable

pce-report sr-te disable

Parameters

rsvp-te {enable | disable}

specifies to enable or disable PCE reporting for all RSVP-TE LSPs

sr-te {enable | disable}

specifies to enable or disable PCE reporting for all SR-TE LSPs

resignal-timer

Syntax

resignal-timer *minutes*

no resignal-timer

Context

config>router>mpls

Description

This command specifies the value for the LSP resignal timer. The resignal timer is the time, in minutes, that the 7705 SAR software waits before attempting to resignal the LSPs.

When the resignal timer expires, if the newly computed path for an LSP has a better metric than that for the currently recorded hop list, an attempt is made to resignal that LSP using the make-before-break (MBB) mechanism. If the attempt to resignal an LSP fails, the LSP will continue to use the existing path and a resignal will be attempted the next time the timer expires.

When the resignal timer expires, a trap and syslog message are generated.

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

Default

no resignal-timer

Parameters

minutes

specifies the time the software waits before attempting to resignal the LSPs, in minutes

Values 30 to 10080

srlg-frr

Syntax

srlg-frr [**strict**]

no srlg-frr

Context

config>router>mpls

Description

This system-wide command enables or disables 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 the system. When **srlg-frr** is enabled, CSPF includes the SRLG constraint in the computation of an FRR bypass or detour LSP for protecting the primary LSP path.

The **strict** option is a system-wide option that forces the CSPF to consider any configured SRLG membership lists in its calculation of every LSP path.

CSPF prunes all links with interfaces that belong to the same SRLG as the interface 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-TE 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 set up and the MPLS/RSVP-TE task keeps 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, then the FRR bypass or detour LSP is set up.

An FRR bypass or detour LSP 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-TE task is searching for an SRLG bypass tunnel to associate with the primary path of the protected LSP, the task does the following steps.

- First, the task checks for any configured manual bypass LSP that has CSPF enabled and that satisfies the SRLG constraints.
- The task skips any non-CSPF bypass LSP since there is no ERO returned with which to check the SRLG constraint.
- If no path is found, the task checks for an existing dynamic bypass LSP that satisfies the SRLG and other primary path constraints.
- If no bypass path is found, then the task makes a request to CSPF to try to create one.

Once the primary path of the LSP is set up and is operationally up, any subsequent changes to the SRLG membership of an interface that the primary path is using will not be considered by the MPLS/RSVP-TE 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 due to a failure or to a make-before-break (MBB) operation. A make-before-break operation occurs as a result of a global revertive operation, a reoptimization of the LSP path (timer-based or manual), or a change by the user to any of the path constraints.

Once the FRR bypass or detour LSP is set up and is operationally up, any subsequent change to the SRLG membership of an interface that the FRR bypass or detour LSP is using will not be considered by

the MPLS/RSVP-TE task at the PLR until the next opportunity that the association with the primary LSP path is rechecked. The association is rechecked if the FRR bypass or detour LSP is reoptimized. Detour routes are not reoptimized and are resignaled if the primary path is down.

The user must first shut down MPLS before enabling or disabling the **srlg-frr** option in CLI.

An RSVP-TE interface can belong to multiple SRLGs. The user creates SRLGs using the **config>router>mpls>srlg-group** command. The user associates the SRLGs with an RSVP-TE interface using the **srlg-group** command in the **config>router>mpls>interface** context.

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

Default

no srlg-frr

Parameters

strict

specifies that the CSPF calculation for the FRR backup must include the SRLG constraint and the backup must be on the resulting list of eligible backup paths

Values non-strict: srlg-frr
 strict: srlg-frr **strict**

sr-te-resignal-timer

Syntax

sr-te-resignal-timer *minutes*

no sr-te-resignal-timer

Context

config>router>mpls

Description

This command specifies the value for the SR-TE LSP resignal timer when the path computation method is set to the local CSPF or the PCE algorithm.

The resignal timer is the time, in minutes, that the 7705 SAR software waits before attempting to reoptimize all paths of all SR-TE LSPs. The reoptimization is performed by either the local CSPF or the PCE algorithm, as configured by the **path-computation-method** command.

When local CSPF is used and the resignal timer expires, the 7705 SAR provides the current path of the SR-TE LSP and the TE-DB updates the total IGP or TE metric of the current path and checks the validity of the hops and labels. CSPF then computes a new path for each SR-TE LSP. MPLS programs the new path only if the total metric of the new computed path is different from the updated metric of the current path, or if one or more hops or labels of the current path are invalid. Otherwise, the current path is considered to be one of the most optimal ECMP paths and is not updated in the datapath.

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

Default

no sr-te-resignal-timer

Parameters

minutes

specifies the time the software waits before attempting to resignal the SR-TE LSPs

Values 30 to 10080

3.24.2.1.3 MPLS-labels commands

mpls-labels

Syntax

mpls-labels

Context

config>router

Description

This command enables the context to configure MPLS label parameters.

reserved-label-block

Syntax

[no] reserved-label-block *name*

Context

config>router>mpls-labels

Description

This command configures a label block for a range of labels that are reserved for local assignment to specific applications, such as segment routing (SR) adjacency SIDs. The labels are reserved from the dynamic label range. The reserved label block is not advertised by the IGP. Up to four reserved label blocks can be configured on a system. Only one range can be configured for each block.

The **no** form of this command removes the reserved label block.

Default

none

Parameters

name

the name of the reserved label block, up to 64 characters

start-label

Syntax

start-label *start-value* **end-label** *end-value*

no start-label

Context

config>router>mpls-labels>reserved-label-block

Description

This command configures start and end labels for the reserved label block. This command must be configured in order for the reserved label block to be created.

Default

start-label 0, end-label 0

Parameters

start-value

specifies a starting value for the label range

Values 18432 to 524287 (within the dynamic label range)

end-value

specifies an ending value for the label range

Values 18432 to 524287 (within the dynamic label range)

sr-labels

Syntax

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

no sr-labels

Context

config>router>mpls-labels

Description

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

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

Default

no sr-labels

Parameters

start-value

specifies the start label value in the SRGB

Values 18432 to 524287 within dynamic label range

end-value

specifies the end label value in the SRGB

Values 18432 to 524287 within dynamic label range

static-label-range

Syntax

static-label-range *static-range*

no static-label-range

Context

config>router>mpls-labels

Description

This command configures the range of MPLS static label values shared among static LSP, MPLS-TP LSP, and static service VC labels. When this range is configured, it is reserved and cannot be used by other protocols such as RSVP-TE, LDP, BGP, or segment routing to assign a label dynamically.

Default

static-label-range

Parameters

static-range

specifies the size of the static label range in number of labels. The minimum label value in the range is 32. The maximum label value is computed as $\{32 + \textit{static-range} - 1\}$.

Values 0 to 131040

Default 18400

3.24.2.1.4 RSVP LSP statistics commands

ingress-statistics

Syntax

ingress-statistics

Context

config>router>mpls

Description

This command enters the context to configure LSP ingress statistics.

Default

n/a

lsp

Syntax

[no] lsp *lsp-name* **sender** *ip-address*

Context

config>router>mpls>ingress-statistics

Description

This command configures statistics in the ingress data path of a terminating RSVP LSP at an egress LER. The LSP name must correspond to the name configured by the user at the ingress LER. It must not contain a colon (:), which is used as a field separator by the ingress LER for encoding the LSP and path names into the RSVP Session Name field in the Session_Attribute object. The user must also execute the **no shutdown** command in this context to enable statistics collection.

The **no** form of this command disables statistics for this RSVP LSP in the ingress data path and removes the accounting policy association from the LSP.

Default

n/a

Parameters

lsp-name

the LSP name as configured at the ingress LER, up to 32 characters in length

ip-address

the IP address of the ingress LER for the LSP

accounting-policy

Syntax

accounting-policy *policy-id*

no accounting-policy

Context

config>router>mpls>ingr-stats>lsp

config>router>mpls>lsp>egress-statistics

Description

This command associates an accounting policy with an RSVP LSP. Only one accounting policy at a time can be associated with an RSVP LSP on a particular node.

An accounting policy must first be configured in the **config>log>accounting-policy** context before it can be associated; otherwise an error message is generated.

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

Default

no accounting-policy

Parameters

policy-id

the accounting policy ID

Values 1 to 99

collect-stats

Syntax

[no] collect-stats

Context

config>router>mpls>ingr-stats>lsp

config>router>mpls>lsp>egress-statistics

Description

This command enables accounting and statistical data collection.

The collected statistic counters can be retrieved via **show** and **monitor** commands or with the SNMPv3 interface. The counters can be saved to an accounting file if the specific statistics collection record is included in the default accounting policy or in a user-defined accounting policy.

If the **no collect-stats** command is issued, the statistics are still accumulated by the forwarding engine. However, the CPU will not obtain the results and write them to the accounting file. If a subsequent **collect-stats** command is issued, then the counters written to the accounting file will include all the traffic that went through while the **no collect-stats** command was in effect.

Default

no collect-stats

egress-statistics

Syntax

[no] egress-statistics

Context

config>router>mpls>lsp

Description

This command configures statistics in the egress data path of an originating LSP at a head-end node. The user must also execute the **no shutdown** command in this context to enable statistics collection.

The **no** form of this command disables the statistics in the egress data path and removes the accounting policy association from the RSVP LSP.

Default

no egress-statistics

3.24.2.1.5 Interface commands

interface

Syntax

[no] interface *ip-int-name*

Context

config>router>mpls

Description

This command enables MPLS protocol support on an IP interface. MPLS commands are not executed on an IP interface where MPLS is not enabled.

The **no** form of this command deletes all MPLS commands that are defined under the interface, such as **label-map**. The interface must be shut down before it can be deleted. If the interface is not shut down, the **no interface *ip-int-name*** command issues a warning message on the console indicating that the interface is administratively up.

Default

shutdown

Parameters

ip-int-name

identifies the network IP interface. The interface name character string cannot be in the form of an IP address. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), 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

Description

This command associates admin groups with this interface. The admin group must already be defined in the **config>router>if-attribute** context (see the 7705 SAR Router Configuration Guide, "IP Router Command Reference").

Up to five groups can be specified with one command. When an admin group is bound to one or more interfaces, its value cannot be changed until all bindings are removed.

When associated with MPLS interfaces, the interfaces can be included or excluded in the LSP path definition by matching on the admin-group name. CSPF will calculate a path that satisfies the admin-group include and exclude constraints.

The configured admin-group membership is applied in all levels or areas that the interface is participating in. The same interface cannot have different memberships in different levels or areas.

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

The **no** form of this command deletes the association of this interface with one or more of the admin groups.

Default

no admin-group

Parameters

group-name

specifies the name of the group. The group names should be the same across all routers in the MPLS domain.

srlg-group

Syntax

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

Context

config>router>mpls>interface

Description

This command associates an MPLS interface with one or more SRLGs. The SRLG must already be defined in the **config>router>if-attribute** context (see the 7705 SAR Router Configuration Guide, "IP Router Command Reference").

Up to five SRLGs can be specified with one command, and the command can be issued multiple times. When an SRLG is bound to one or more interfaces, its value cannot be changed until all bindings are removed.

When SRLGs are associated with MPLS interfaces, CSPF at an LER will exclude the SRLGs of interfaces used by the LSP primary path when calculating the route of the secondary path. CSPF at an LER or LSR will also exclude the SRLGs of the outgoing interface of the primary LSP path in the calculation of the path of the FRR backup LSP. This provides a path disjoint between the primary path and the secondary path or FRR backup path of an LSP.

The configured SRLG membership is applied in all levels or areas that the interface is participating in. The same interface cannot have different memberships in different levels or areas.

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

The **no** form of this command deletes the association of this interface with one or more of the SRLGs.

Default

n/a

Parameters

group-name

specifies the name of the SRLG. The group names should be the same across all routers in the MPLS domain.

te-metric

Syntax

te-metric *value*

no te-metric

Context

config>router>mpls>interface

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 OSPF-TE metric is encoded as a sub-TLV type 5 in the Link TLV. The metric value is encoded as a 32-bit unsigned integer. The IS-IS-TE metric is encoded as sub-TLV type 18 as part of the extended IS reachability TLV. The metric value is encoded as a 24-bit unsigned integer.

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. Such constraints include bandwidth, admin-groups, and hop limit. Then, CSPF will 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 in CSPF LSP path computation can be configured by entering the command **config>router>mpls>lsp *lsp-name*>metric-type te** or **config>router>mpls>lsp-template *template-name*>cspf use-te-metric**.

The TE metric is only used in CSPF computations for MPLS paths and not in the regular SPF computation for IP reachability.

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

Default

no te-metric

Parameters

value

1 to 16777215

3.24.2.1.6 Interface label-map commands

label-map

Syntax

[no] label-map *in-label*

Context

config>router>mpls>interface

Description

This command is used on either transit or egress LSP routers when a static LSP is defined. The static LSP on the ingress router is initiated using the **config>router>mpls>static-lsp** *lsp-name* command. The *in-label* is associated with a **pop** action 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

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. Once 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*.

Default

n/a

swap

Syntax

swap *out-label* **nexthop** *ip-address*

no swap

Context

config>router>mpls>if>label-map

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*.

Default

n/a

Parameters

out-label

specifies the label value to be swapped with the *in-label*. Label values 16 through 1048575 are defined as follows:

Label values 16 through 31 are 7705 SAR reserved

Label values 32 through 1023 are available for static assignment

Label values 1024 through 2047 are reserved for future use

Label values 2048 through 18431 are statically assigned for services

Label values 28672 through 131071 are dynamically assigned for both MPLS and services

Label values 131072 through 1048575 are reserved for future use

Values 16 to 1048575

ip-address

specifies the IP address to forward to. If an ARP entry for the next hop exists, then the static LSP will be marked operational. If an ARP entry does not exist, software will set the operational status of the static LSP to down and continue to ARP for the configured next-hop at a fixed interval.

3.24.2.1.7 LSP and LSP template commands

lsp

Syntax

[no] **lsp** *lsp-name* [**bypass-only** | **sr-te**]

Context

config>router>mpls

Description

This command creates an LSP that is signaled dynamically by the 7705 SAR.

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.

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 shut down and unbound from all SDPs before it can be deleted.

Default

n/a

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

defines 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 7705 SAR selects it. If no manual bypass tunnel is found, the 7705 SAR dynamically signals a bypass LSP as the default behavior. The CLI for this feature includes a command that provides the user with the option to disable dynamic bypass creation on a per-node basis.

sr-te

defines an LSP as a segment routing LSP exclusively

lsp-template

Syntax

lsp-template *template-name*

lsp-template *template-name* **mesh-p2p**

lsp-template *template-name* **one-hop-p2p**

no lsp-template *template-name*

Context

config>router>mpls

Description

This command creates an LSP template that is referenced when dynamic LSP creation is required. The LSP template type, **mesh-p2p** or **one-hop-p2p**, must be specified when the template is first created.

The **no** form of this command deletes the LSP template. The LSP template cannot be deleted if a client application is using it.

Parameters

template-name

specifies an LSP template name up to 32 characters in length. The LSP template name and the LSP name cannot be the same.

mesh-p2p | one-hop-p2p

specifies the type of LSP the template signals

adaptive

Syntax

[no] adaptive

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Description

This command enables the make-before-break (MBB) functionality for an LSP, LSP path, or LSP instance created using an LSP template. When enabled for the LSP, a make-before-break operation will be performed for the primary path and all the secondary paths of the LSP.

Default

adaptive

adspec

Syntax

[no] adspec

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Description

When enabled, the advertised data (ADSPEC) object will be included in RSVP-TE messages.

Default

no adspec

bfd

Syntax

bfd

Context

config>router>mpls>lsp

config>router>mpls>lsp>primary

```
config>router>mpls>lsp>secondary
```

Description

This command enables the context to configure seamless BFD (S-BFD) parameters on an LSP that is defined as an SR-TE LSP. The configuration in the **config>router>mpls>lsp>secondary** context is valid only for SR-TE LSPs.

bfd-enable

Syntax

[no] bfd-enable

Context

```
config>router>mpls>lsp>bfd
```

```
config>router>mpls>lsp>primary>bfd
```

```
config>router>mpls>lsp>secondary>bfd
```

Description

This command enables S-BFD for an SR-TE LSP. The BFD template must be configured in the **config>router>bfd** context before issuing the **bfd-enable** command.

The **no** form of this command disables S-BFD on the SR-TE LSP.

Default

no bfd-enable

bfd-template

Syntax

bfd-template *name*

no bfd-template

Context

```
config>router>mpls>lsp>bfd
```

```
config>router>mpls>lsp>primary>bfd
```

```
config>router>mpls>lsp>secondary>bfd
```

Description

This command specifies the BFD template to be used for an S-BFD session on an SR-TE LSP. The BFD template is configured in the **config>router>bfd** context and defines the parameters used during the session, such as the minimum transmit and receive intervals for control packets.

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

Default

no bfd-template

Parameters

name

the name of the existing BFD template

failure-action

Syntax

failure-action failover-or-down

no failure-action

Context

config>router>mpls>lsp>bfd

Description

This command configures the resulting action when BFD fails on an SR-TE LSP.

A failure action of **failover-or-down** means that a switchover from the active path is triggered if the BFD session fails on the active path (primary or standby). If there is no available path to switch to, the LSP is taken operationally down. This failure action is only applicable to SR-TE LSPs.

The system generates an SNMP trap if BFD fails on an LSP, regardless of whether a failure action is configured.

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

Default

no failure-action

Parameters

failover-or-down

the active path of an SR-TE LSP switches to the secondary path, or the LSP goes operationally down if no other path is available

return-path-label

Syntax

return-path-label *label-value*

no return-path-label

Context

config>router>mpls>lsp>bfd

```
config>router>mpls>lsp>primary>bfd
config>router>mpls>lsp>secondary>bfd
```

Description

This command configures the Seamless Bidirectional Forwarding Detection (S-BFD) session to echo mode and adds an additional MPLS label, referring to an MPLS-labeled reply path for the S-BFD packet, to the bottom of the label stack for the S-BFD packet. The command applies to the initiator of the S-BFD sessions. The return-path label may be a binding SID or other MPLS path configured on the reflector router. Instead of being routed through the IGP path, the S-BFD packet returns to the initiator through this MPLS return path.

The **no** form of this command disables the controlled return-path label and echo mode for S-BFD. S-BFD returns to asynchronous mode and the initiator node does not push a return-path label. Any S-BFD packets for this LSP or path that the reflector receives are sent back using a routed return path.

Default

disabled

Parameters

label-value

Specifies the label value

Value	32 to 524287
--------------	--------------

wait-for-up-timer

Syntax

wait-for-up-timer *seconds*

no wait-for-up-timer

Context

```
config>router>mpls>lsp>bfd
config>router>mpls>lsp>primary>bfd
config>router>mpls>lsp>secondary>bfd
```

Description

This command configures the time to wait for BFD to come up on a path. This timer applies to SR-TE LSPs. The timer is started when BFD is first enabled on a path or BFD transitions from up to down. If the timer expires before BFD comes up, the path is torn down and the retry timer is started.

The **wait-for-up-timer** command is only available if the configured failure action is **failover-or-down**.

The **no** form of the command resets the timer to the default value.

Default

4

Parameters

seconds

the time to wait for BFD to come up

Values 0 to 60

bgp-transport-tunnel

Syntax

bgp-transport-tunnel {**include** | **exclude**}

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Description

This command allows an RSVP-TE LSP to be used as a transport LSP for BGP tunnel routes or blocks it from being used.

Default

include

Parameters

include

allows an RSVP-TE LSP to be used as a transport LSP from the ASBR to a local PE router, from an ingress PE to the ASBR in the local AS or between multihop EBGP peers with ASBR-to-ASBR adjacency

exclude

blocks an RSVP-TE LSP from being used as a transport LSP from the ASBR to a local PE router, from an ingress PE to the ASBR in the local AS or between multihop EBGP peers with ASBR-to-ASBR adjacency

cspf

Syntax

cspf [**use-te-metric**]

no cspf

Context

config>router>mpls>lsp-template

Description

This command enables constrained shortest path first (CSPF) computation for constrained-path LSPs. Constrained-path LSPs are the LSPs that take configuration constraints into account. CSPF is also used to calculate the FRR bypass or detour LSP routes when fast reroute is enabled.

When an LSP template is created, CSPF is automatically enabled and cannot be disabled. The metric is set to IGP by default but can be changed to TE with the **use-te-metric** parameter.

Default

cspf

Parameters

use-te-metric

specifies to use the TE metric for the purpose of LSP path computation by CSPF

default-path

Syntax

default-path *path-name*

Context

config>router>mpls>lsp-template

Description

This command specifies the default path to be used for signaling an LSP created using the LSP template. You must reference a default path before the template is put in a **no shutdown** state.

Parameters

path-name

specifies the default path name to be used

entropy-label

Syntax

entropy-label {**force-disable** | **inherit** | **enable**}

Context

config>router>mpls>lsp

Description

This command configures the use of entropy labels for an LSP.

If **entropy-label** is enabled, the entropy label and entropy label indicator (ELI) are inserted in the label stack. In some cases, this may result in an unsupported label stack depth or large changes in the label

stack depth during the lifetime of an LSP (for example, due to switching from a primary path with ELC enabled to a secondary path for which the far end has not signaled ELC).

This command provides local control at the head end of an RSVP LSP over whether an entropy label is inserted on an LSP by overriding the ELC signaled from the far-end LER, and control over how the additional label stack depth is accounted for.

By default, the value of **entropy-label** is inherited from the MPLS level. This command overrides the default MPLS behavior on a per-LSP basis. For auto-LSPs, it can only be configured in LSP templates of type one-hop-p2p and mesh-p2p.

Under the LSP context, when the value of **entropy-label** is set to **enable**, the ingress LER considers what is signaled from the egress node for ELC when marking the NHLFE as entropy-label-capable. When the value of **entropy-label** is set to **enable** at the LSP level, the system always marks the LSP as entropy label-capable regardless of the signaled value, to ensure that the potential additional label stack depth is accounted for. In this scenario, the TTM and NHLFE can be out of synchronization based on what is configured at the egress node. That is, the application will always account for the entropy label and ELI in the label stack without taking into consideration the signaled value of the ELC.

When the value of **entropy-label** changes at either the MPLS level or the LSP level, the new operational value does not take effect until the LSP is resignaled. A **shutdown/no shutdown** command must be performed to enable the new value.

The user can use the **clear** command or bounce MPLS using the **shutdown/no shutdown** command to force the new value to take effect for a large numbers of LSPs.

Default

entropy-label inherit

Parameters

force-disable

the ingress LER will not consider the entropy label and ELI in the label stack while sending the information to the TTM and NHLFE. The system will mark the TTM and NHLFE as ELC not supported, and applications will not insert an entropy label or entropy label indicator in the label stack.

inherit

the value of **entropy-label** is inherited from the setting in the MPLS context

enable

the ingress LER will take into consideration what is signaled from the egress node for ELC for marking the NHLFE, while the TTM is always marked. Although applications will only insert the entropy label if the far end signals ELC, the additional two labels of the entropy label and ELI are always accounted for.

exclude

Syntax

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

Context

```
config>router>mpls>lsp
config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary
config>router>mpls>lsp>template
```

Description

This command specifies the admin groups to be excluded when an LSP is set up. Up to 5 groups per operation can be specified, up to 32 maximum. The admin groups are defined in the **config>router>if-attribute** context.

Use the **no** form of the command to remove the exclude command.

Default

no exclude

Parameters

group-name

specifies the existing group name to be excluded when an LSP is set up

fallback-path-computation-method

Syntax

fallback-path-computation-method [**local-cspf** | **none**]

no fallback-path-computation-method

Context

```
config>router>mpls>lsp
```

Description

This command specifies the fallback path computation method used if all configured PCEs are down or are signaling overload and the redelegation timer has expired. This method is used regardless of whether the LSP is PCE-controlled and PCE-computed, or just PCE-computed.

The **no** form of this command removes the fallback path computation method used.

Default

fallback-path-computation-method none

Parameters

local-cspf

specifies to fall back to using local CSPF computation

none

specifies to fall back to using the hop-to-label computation for SR-TE LSPs

fast-reroute

Syntax

[no] **fast-reroute** [*frr-method*]

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Description

This command creates a precomputed protection LSP from each node in the path of the LSP. If a link or LSP failure occurs between two nodes, traffic is immediately rerouted on the precomputed protection LSP. When **fast-reroute** is specified, the default **fast-reroute** method is the one-to-one method.

When **fast-reroute** is enabled, each node along the path of the LSP tries to establish a protection LSP as follows.

- Each upstream node sets up a protection LSP that avoids only the immediate downstream node, and merges back onto the actual path of the LSP as soon as possible.
- If it is not possible to set up a protection LSP that avoids the immediate downstream node, a protection LSP can be set up to the downstream node on a different interface.
- The protection LSP may take one or more hops (see [igp-shortcut](#)) before merging back onto 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.

FRR 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-TE to set up their protection LSP. TE must be enabled for fast reroute to work.

CSPF must be enabled for fast reroute to work. If an LSP is configured with **fast-reroute** *frr-method* specified but does not enable CSPF, neither global revertive nor local revertive will be available for the LSP to recover.

The one-to-one fast reroute method creates a separate detour LSP for each backed-up LSP. One-to-one is not supported for LSP templates.

The facility fast reroute method, sometimes called many-to-one, takes advantage of the MPLS label stack. Instead of creating a separate LSP for every backed-up LSP, a single LSP is created that serves to back up a set of LSPs. This LSP tunnel is called a bypass tunnel. The bypass tunnel must intersect the path of the original LSPs somewhere downstream of the point of local repair (PLR). This constrains the set of LSPs being backed up via that bypass tunnel to those LSPs that pass through a common downstream node. All LSPs that pass through the PLR and through this common node that do not also use the facilities involved in the bypass tunnel are candidates for this set of LSPs.

The **no** form of the **fast-reroute** command removes the protection 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.

Default

no fast-reroute

Parameters

frr-method

specifies the fast reroute method to use

Values one-to-one, facility

Default one-to-one

hop-limit

Syntax

hop-limit *limit*

no hop-limit

Context

config>router>mpls>lsp>fast-reroute

config>router>mpls>lsp-template>fast-reroute

Description

For fast reroute, this command defines how many more routers a protection tunnel is allowed to traverse compared with the LSP itself. For example, if an LSP traverses four routers, any protection tunnel for the LSP can be no more than 10 router hops, including the ingress and egress routers.

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

Default

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  
config>router>mpls>lsp-template>fast-reroute
```

Description

This command enables or disables 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.

When **node-protect** is enabled, the 7705 SAR provides node protection on the specified LSP. If node protection cannot be provided, link protection is attempted. If link protection cannot be provided, there will be no protection.

The **no** form of this command provides link protection. If link protection cannot be provided, there will be no protection.

Default

node-protect (for an LSP)
no node-protect (for an LSP template)

propagate-admin-group

Syntax

[no] propagate-admin-group

Context

```
config>router>mpls>lsp>fast-reroute  
config>router>mpls>lsp-template>fast-reroute
```

Description

The command enables the signaling of the primary LSP path admin-group constraints in the FAST_REROUTE object at the ingress LER.

When this command is executed, the admin-group constraints configured in the context of the point-to-point 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.

The ingress LER propagates these constraints to the downstream nodes during the signaling of the LSP so that the downstream nodes can include the constraints in the selection of the FRR backup LSP for the LSP primary path.

The ingress LER inserts the FAST_REROUTE object by default in a primary LSP PATH message.

The same admin-group constraints can be copied into the SESSION_ATTRIBUTE object using the [propagate-admin-group](#) command at the **config>router>mpls>lsp** level. They are intended for the use of an LSR, typically an ABR, to expand the ERO of an inter-area LSP path. They 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 whether the hop is strict or loose.

The primary path admin-group constraints can be copied into the FAST_REROUTE object only, the SESSION_ATTRIBUTE object only, or both. The PLR rules for processing the admin-group constraints can make use of either of the two objects.

If the FAST_REROUTE object is disabled (no [frr-object](#)), the admin-group constraints will not be propagated.

Default

no propagate-admin-group

from

Syntax

from *ip-address*

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Description

This command specifies 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. This command is only allowed for an LSP template of type **mesh-p2p**.

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 due to 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.

Default

system IP address

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-template

Description

This command 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. This value can be changed dynamically for an LSP that is already set up. If the new value is less than the current number of hops of the established LSP, the LSP is brought down. The 7705 SAR 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 of hops of the established LSP, the LSP is not affected.

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

Default

255 (LSP and LSP mesh template)

2 (one-hop template)

Parameters

number

specifies the number of hops the LSP can traverse, expressed as an integer

Values 2 to 255

igp-shortcut

Syntax

igp-shortcut [*lfa-protect* | *lfa-only*] [*relative-metric* [*offset*]]

no igp-shortcut

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Description

This command enables the use of an RSVP-TE LSP by OSPF or IS-IS routing protocols as a shortcut or as a forwarding adjacency for resolving IGP routes.

When the **rsvp-shortcut** or the **advertise-tunnel-link** command is enabled at the OSPF or IS-IS instance level, all RSVP-TE LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured with the **config>router>mpls>lsp>to** command, corresponds to a router ID of a remote node.

If the command is used with no options, and the **rsvp-shortcut** command is enabled, the LSP is included in the main SPF but not in the LFA SPF algorithm. If the **advertise-tunnel-link** command is enabled, the tunnel is advertised as a point-to-point link if it has the best LSP metric, is included in the main SPF if advertised, but is not included in the LFA SPF algorithm.

If the command is used with the **lfa-protect** option, and the **rsvp-shortcut** command is enabled, an LSP can be included in both the main SPF and the LFA SPF algorithm. If the **advertise-tunnel-link** command is enabled, the tunnel is advertised as a point-to-point link if it has the best LSP metric, is included in the main SPF if advertised, and is included in the LFA SPF algorithm whether it is advertised or not.

For a given prefix, the LSP can be used either as a primary next hop or as an LFA next hop, but not both. If the main SPF calculation selects a tunneled primary next hop for a prefix, the LFA SPF calculation will not select an LFA next hop for this prefix and the protection of this prefix will rely on the RSVP-TE LSP FRR protection. If the main SPF calculation selects a direct primary next hop, the LFA SPF calculation will select an LFA next hop for this prefix but will prefer a direct LFA next hop over a tunneled LFA next hop.

If the command is used with the **lfa-only** option, and the **rsvp-shortcut** command is enabled, an LSP can be included in the LFA SPF algorithm only. If the **advertise-tunnel-link** command is enabled, the tunnel is not advertised as a point-to-point link, is not included in the main SPF, and is only included in the LFA SPF algorithm.

When the **lfa-only** option is enabled, the introduction of IGP shortcuts does not affect the main SPF decision. For a given prefix, the main SPF calculation always selects a direct primary next hop. The LFA SPF calculation will select an LFA next hop for this prefix but will prefer a direct LFA next hop over a tunneled LFA next hop.

When the **relative-metric** option is enabled, IGP will apply the shortest IGP cost between the endpoints of the LSP plus the value of the offset (instead of the LSP operational metric) when calculating the cost of a prefix that is resolved to the LSP. The offset value is optional and defaults to zero. The minimum net cost for a prefix is one (1) after applying the offset. The tunnel table manager (TTM) continues to show the LSP operational metric as provided by MPLS; therefore, applications such as BGP and static route shortcuts will continue to use the LSP operational metric.

The **relative-metric** option and the **lfa-protect** or the **lfa-only** options are mutually exclusive. An LSP with the **relative-metric** option enabled cannot be included in the LFA SPF calculation and the **relative-metric** option cannot be enabled if the LSP is included in the LFA SPF calculation when the **rsvp-shortcut** option is enabled in the IGP.

The **relative-metric** option is ignored when forwarding adjacency is enabled in OSPF or IS-IS. In this case, IGP advertises the LSP as a point-to-point unnumbered link along with the LSP operational metric as returned by MPLS and capped to the maximum link metric allowed in that IGP.

Both the main SPF and the LFA SPF algorithms use the local IGP database to resolve the routes.

The **no** form of this command disables the use of an RSVP-TE LSP by OSPF or IS-IS as a shortcut or a forwarding adjacency for resolving IGP routes.

For more information about IGP shortcuts and LFA, see the 7705 SAR Routing Protocols Guide, "LDP and IP fast reroute (FRR) for OSPF prefixes" and "LDP and IP fast reroute (FRR) for IS-IS prefixes".

Default

igp-shortcut – all RSVP-TE LSPs originating on this node are eligible by default as long as the destination address of the LSP corresponds to a router ID of a remote node

Parameters

lfa-protect

an LSP is included in both the main SPF and the LFA SPF calculation

lfa-only

an LSP is included in the LFA SPF calculation only

relative-metric [offset]

the shortest IGP cost between the endpoints of the LSP plus the configured offset, instead of the LSP operational metric returned by MPLS, is used when calculating the cost of prefix resolved to this LSP. The *offset* parameter is optional. If the **relative-metric** option is enabled without specifying the *offset* parameter value, a value of 0 is used.

Values –10 to +10

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

Description

This command specifies the admin groups to be included when an LSP is set up. Up to 5 groups per operation can be specified, and up to 32 maximum.

The **no** form of the command deletes the specified groups in the specified context.

Default

no include

Parameters

group-name

specifies admin groups to be included when an LSP is set up

label-stack-reduction

Syntax

[no] label-stack-reduction

Context

config>router>mpls>lsp

Description

This command enables label stack size reduction for an SR-TE LSP.

The label stack reduction algorithm attempts to replace adjacency SIDs corresponding to a segment of the explicit path with a node SID whenever the constraints of the path are met by all the ECMP paths to that node SID.

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

Default

no label-stack-reduction

least-fill

Syntax

[no] least-fill

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Description

This command enables the use of the least-fill path selection method for the computation of the path of this LSP.

When MPLS requests the computation of a path for this LSP, CSPF finds all equal-cost shortest paths that satisfy the constraints of this path. Then, CSPF identifies the single link in each of these paths that has the least available bandwidth as a percentage of its maximum reservable bandwidth. It then selects the path that has the highest percentage available bandwidth. CSPF identifies the least-available bandwidth link in each equal-cost path after it has accounted for the bandwidth of the new requested path of this LSP.

CSPF applies the least-fill path selection method to all requests for a path, primary and secondary, of an LSP for which this option is enabled. The bandwidth of the path can be any value, including zero.

MPLS resignals and moves the LSP to the new path in the following cases:

- initial LSP path signaling
- retry of an LSP path after failure

- MBB due to an LSP path configuration change, that is, a user change to the bandwidth parameter of the primary or secondary path, or a user enabling of the fast-reroute option for the LSP
- MBB of the path due to an update to the primary path SRLG
- MBB due to fast reroute global revertive procedures on the primary path
- manual resignaling of an LSP path or of all LSP paths by the user

During a manual resignaling of an LSP path, MPLS always resignals the path even if the new path is the same as the current path and even if the metric of the new path is the same as the metric of the current path.

During a timer-based resignaling of an LSP path that has the least-fill option enabled, MPLS only resignals the path if the metric of the new path is different from the metric of the current path.

Default

no least-fill - the path of an LSP is randomly chosen among a set of equal-cost paths

load-balancing-weight

Syntax

load-balancing-weight *weight*

no load-balancing-weight

Context

config>router>mpls>lsp>load-balancing-weight

config>router>mpls>lsp-template>load-balancing-weight

Description

This command assigns a weight to an MPLS LSP or an MPLS LSP template for use in weighted load balancing (weighted ECMP) over MPLS.

The **no** form of this command removes the load-balancing weight from the LSP or LSP template.

Default

0

Parameters

number

a 32-bit integer representing the weight of the LSP

Values 1 to 100

local-sr-protection

Syntax

local-sr-protection *local-sr-protection*

no local-sr-protection

Context

config>router>mpls>lsp

Description

This command configures the SR LFA protection needed for the adjacencies used in the path computation of an SR-TE LSP by the local CSPF.

The default value of the command is **preferred**. The local CSPF will prefer a protected adjacency over an unprotected adjacency whenever both exist for a TE link. However, the entire computed path can combine both types of adjacencies.

When local SR protection is configured as **mandatory**, the local CSPF uses it as an additional path constraint and selects protected adjacencies exclusively in computing the path of the SR-TE LSP. CSPF returns no path if all candidate paths that otherwise satisfy all other LSP path constraints do not have an unprotected SID for each of their TE links.

When local SR protection is configured as **none**, the local CSPF uses it as an additional path constraint and selects unprotected adjacencies exclusively in computing the path of the SR-TE LSP. CSPF returns no path if all candidate paths that otherwise satisfy all other LSP path constraints do not have a protected SID for each of their TE links.

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

Default

preferred

Parameters

local-sr-protection

the local CSPF protection method for SR-TE LSPs

Values preferred, mandatory, none

max-sr-labels

Syntax

max-sr-labels *label-stack-size*[**additional-frr-labels** *labels*]

no max-sr-labels

Context

config>router>mpls>lsp

Description

This command configures the maximum number of labels that the ingress LER can push for an SR-TE LSP.

This command is used to allow room to insert additional transport, service, and other labels when packets are forwarded in a context.

The **max-sr-labels** *label-stack-size* value should reflect the desired maximum label stack of the primary path of the SR-TE LSP.

The value in **additional-frr-labels** *labels* should reflect additional labels inserted by remote LFA for the backup next hop of the SR-TE LSP.

The sum of both label values represents the worst-case transport of SR label stack size for this SR-TE LSP. The sum is populated by MPLS in the Tunnel Table Manager (TTM) so that services and shortcut applications can check the TTM to determine whether a service can be bound or a route can be resolved to this SR-TE LSP.

The maximum label stack supported by the router is always signaled by the PCC in the PCEP Open object as part of the SR-PCE-CAPABILITY TLV. The maximum label stack is referred to as the Maximum Stack Depth (MSD).

In addition, the per-LSP value for the **max-sr-labels** option, if configured, is signaled by the PCC to the PCE in the Segment-ID (SID) Depth value in a METRIC object for both a PCE-computed LSP and a PCE-controlled LSP. The PCE will compute and provide the full explicit path with TE links specified. If there is no path with the number of hops lower than the MSD value or the SID Depth value (if signaled), a reply with no path will be returned to the PCC.

For a PCC-controlled LSP, if the label stack returned by the TE-DB hop-to-label translation exceeds the per-LSP maximum label stack size for the SR, the LSP is brought down.

Default

max-sr-labels 6

additional-frr-labels 1

Parameters

label-stack-size

specifies the label stack size of the primary path of the SR-TE LSP

Values 1 to 11

Default 6

additional-frr-labels *labels*

sets the number of additional labels inserted by remote LFA for the backup next hop of the SR-TE LSP

Values 0 to 4

Default 1

metric

Syntax

metric *metric*

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Description

This command specifies the metric for this LSP, which is used to select an LSP from among a set of LSPs that are destined for the same egress router. The LSP with the lowest metric will be selected.

The operational metric for an LSP that uses the TE metric in CSPF path calculations can be overridden by the configured administrative LSP metric parameter.

Default

1

Parameters

metric

specifies the metric for this LSP

Values 1 to 16777215

metric-type

Syntax

metric-type *metric-type*

no metric-type

Context

config>router>mpls>lsp

Description

This command specifies the link metric type to use in the SR-TE LSP path computation by either the local CSPF or the PCE algorithm.

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

Default

igp

Parameters

metric-type

specifies the metric for this LSP

Values igp, te

path-computation-method

Syntax

path-computation-method *path-computation-method*

no path-computation-method

Context

config>router>mpls>lsp

Description

This command configures the path computation method of an SR-TE LSP as either local CSPF or the PCE.

If no specific path computation method is configured for an SR-TE LSP, it uses the hop-to-label path computation method.

The **no** form of this command returns to the default path computation method for the LSP.

Default

no path-computation-method

Parameters

path-computation-method

specifies the path computation method for this LSP

Values local-cspf, pce

path-profile

Syntax

path-profile *profile-id* [**path-group** *group-id*]

no path-profile *profile-id*

Context

config>router>mpls>lsp

Description

This command configures the PCE path profile and path group ID.

The PCE supports the computation of disjoint paths for two different LSPs originating or terminating on the same or different PE routers. To indicate this constraint to the PCE, the user must configure the PCE path profile ID and path group ID that the PCE-computed or PCE-controlled LSP belongs to. These parameters are passed transparently by the PCC to the PCE and are therefore opaque data to the router.

The association of the optional path group ID is to allow the PCE to determine which profile ID this path group ID must be used with. One path group ID is allowed per profile ID. The user can, however, enter the same path group ID with multiple profile IDs by executing this command multiple times. A maximum of five entries of **path-profile [path-group]** 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-associations

Syntax

pce-associations

Context

config>router>mpls>lsp

Description

This command enables the context to configure LSP bindings with one or more PCEP association groups.

diversity

Syntax

[no] diversity *diversity-assoc-name*

Context

config>router>mpls>lsp>pce-assoc

Description

This command binds the LSP to the specified diversity association group. The diversity association name must exist under the PCC. Up to five diversity associations can be configured per LSP.

The **no** form of the command removes the LSP binding from the specified diversity association group.

Parameters

diversity-assoc-name
the name of an existing diversity association

policy

Syntax

[no] **policy** *policy-assoc-name*

Context

config>router>mpls>lsp>pce-assoc

Description

This command binds the LSP to the specified policy association group. The policy association name must exist under the PCC. Up to five policy associations can be configured per LSP.

The **no** form of the command removes the LSP binding from the specified policy association group.

Parameters

policy-assoc-name
the name of an existing policy association

pce-control

Syntax

[no] **pce-control**

Context

config>router>mpls>lsp

Description

This command enables a PCE-controlled LSP mode of operation for RSVP-TE LSPs and SR-TE LSPs.

The **pce-control** option means that the PCC router delegates full control of the LSP to the PCE (PCE-controlled). Enabling PCE control means that the PCE is acting in active stateful mode for this LSP; the PCE will be able to reroute the path following a failure or reoptimize the path and update the router without the PCC router requesting the update.

The user can delegate CSPF and non-CSPF LSPs or LSPs that have the **path-computation-method pce** option enabled or disabled. The LSP maintains its latest active path computed by the PCE or the PCC router at the time it is delegated. The PCE will only make an update to the path at the next network event or reoptimization.

If PCE reporting is disabled for the LSP, either due to inheritance from the MPLS-level configuration or due to 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

config>router>mpls>lsp-template

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.

The global MPLS-level **pce-report** command can be used to enable or disable PCE reporting for all RSVP-TE LSPs during PCE LSP database synchronization (see **config>router>mpls>pce-report**).

The LSP-level **pce-report** command overrides the global configuration for the reporting of 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

enables PCE reporting

disable

disables PCE reporting

inherit

inherits the global configuration for PCE reporting

propagate-admin-group

Syntax

[no] **propagate-admin-group**

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Description

This command enables propagation of the SESSION_ATTRIBUTE object with resource affinity (C-type 1) in the PATH message. If a SESSION_ATTRIBUTE object with resource affinity is received at an LSR, the LSR will check the compatibility of admin groups received in the PATH message against 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** option enabled and the message includes admin groups, then 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 this 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.

If the configuration of this command is changed (enabled or disabled), the LSP will perform a make-before-break (MBB).

Default

no propagate-admin-group

retry-limit

Syntax

retry-limit *number*

no retry-limit

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Description

This optional command specifies the number of attempts software should make to re-establish the LSP after it has failed. After each successful attempt, the counter is reset to zero.

When the specified number is reached, no more attempts are made and the LSP path is put into the **shutdown** state.

Use the **config>router>mpls>lsp /lsp-name>no shutdown** command to bring up the path after the retry limit is exceeded.

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

Default

0

Parameters

number

specifies the number of times that the 7705 SAR software will attempt to re-establish the LSP after it has failed. Allowed values are integers in the range of 0 to 10000, where 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

Description

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

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

Description

This command specifies the RSVP-TE 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

se

Parameters

ff

fixed filter is 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

shared explicit is 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

config>router>mpls>lsp>primary

config>router>mpls>lsp>secondary

config>router>mpls>lsp-template

Description

This **lsp** form of this command disables the existing LSP, including the primary and any standby secondary paths.

The **primary** and **secondary** forms of this command administratively disable an LSP path and disable an existing LSP. Shutting down an LSP path does not change other configuration parameters for the LSP path.

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 **lsp-template** form of this command disables the existing LSP template.

Use the **no** form of this command to restart the LSP or LSP template. LSPs and LSP templates are created in a shutdown state. Use this command to administratively bring up the LSP or LSP template.

Default

lsp – shutdown

primary – no shutdown

secondary – no shutdown

lsp-template – shutdown

to

Syntax

to *ip-address*

Context

config>router>mpls>lsp

Description

This command specifies the 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.

If the to address does not match the SDP address, the LSP is not included in the SDP definition.

Default

n/a

Parameters

ip-address

specifies the IP address of the egress router

vprn-auto-bind

Syntax

vprn-auto-bind [include | exclude]

Context

config>router>mpls>lsp

config>router>mpls>lsp-template

Description

This command determines whether the associated LSP can be used as part of the autobind feature for VPRN services. By default, an LSP is allowed to be used by the autobind feature.

When VPRN **auto-bind** is set to **exclude**, the associated LSP is not used by the autobind feature for VPRN services.

Default

include

Parameters

include

allows an associated LSP to be used by autobind for VPRN services

exclude

prevents the associated LSP from being used with the autobind feature for VPRN services

3.24.2.1.8 Primary and secondary path commands

primary

Syntax

[no] **primary** *path-name*

Context

config>router>mpls>lsp

Description

This command specifies 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 comparable LSP attributes that are defined in the **config>router>mpls>lsp** context.

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 first in order 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.

Default

n/a

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

Description

This command specifies 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 switchover from the primary path to the secondary path, the 7705 SAR 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 two secondary paths can be specified. Both secondary paths are considered equal, and the first available path is used. The 7705 SAR software will not switch back between secondary paths.

The 7705 SAR software starts signaling 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 in order 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.

Default

n/a

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
```

Description

This command enables the make-before-break (MBB) functionality for an LSP or a primary or secondary LSP path. When enabled for the LSP, a make-before-break operation will be performed for the primary path and all the secondary paths of the LSP.

Default

adaptive

bandwidth

Syntax

```
bandwidth rate-in-mbps  
no bandwidth
```

Context

```
config>router>mpls>lsp>primary  
config>router>mpls>lsp>secondary
```

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).

Default

no bandwidth – bandwidth setting in the global LSP configuration

Parameters

rate-in-mbps

specifies the amount of bandwidth reserved for the LSP path in Mb/s

Values 0 to 100000

hop-limit

Syntax

```
hop-limit number  
no hop-limit
```


Context

```
config>router>mpls>lsp>primary
config>router>mpls>lsp>secondary
```

Description

This optional command overrides the **config>router>mpls>lsp *lsp-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. If the new value is less than the current number of hops of the established LSP, then 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 greater than the current hops of the established LSP, then the LSP will be unaffected.

The **no** form of this command resets the hop limit to the value defined under the LSP definition using the **config>router>mpls>lsp *lsp-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

```
path-preference preference-number
no path-preference
```

Context

```
config>router>mpls>lsp>secondary
```

Description

This command allows a priority value to be assigned to a standby secondary LSP path. The secondary LSP path must be configured in **standby** mode using the **standby** command to ensure that it is signaled and maintained indefinitely in a hot-standby state. The standby secondary LSP path configured with the highest priority (the lowest **path-preference** value) is made the active LSP when the primary LSP is not in use. If multiple standby secondary LSP paths are configured with the same value, the system selects the path with the lowest uptime.

If all standby secondary paths have the default **path-preference** value, a non-standby secondary path remains the active path even though a configured standby secondary path is available.

This command only applies to secondary LSP paths that have been configured in standby mode.

The **no** form of the command resets the **path-preference** *value* to its default.

Default

255

Parameters

preference-number

specifies a priority value for a standby secondary LSP path; the lower the value, the higher the priority.

Values 1 to 255

priority

Syntax

priority *setup-priority*

no **priority**

Context

config>router>mpls>lsp>primary

config>router>mpls>lsp>secondary

Description

This command enables the configuration of LSP setup priority. Nokia routers support path profiling at the tunnel level, where all paths within the same tunnel share the same profile and group value. There can be more than one path in each direction, in which case they all share the same group value. To enforce path diversity or bidirectional behavior for PCE-computed paths, a unique differentiator among LSPs is required. With setup priority, LSPs in each direction can be sorted according to their assigned priority, with 0 as the highest, or "primary", and 7 as the lowest.

The **no** form of this command reverts the LSP path priority to the default values and results in setting the setup priority to 7.

Default

no priority

Parameters

setup-priority

specifies the priority of the reservation for this session at setup time.

Values 0 to 7 (0 is the highest priority and 7 is the lowest priority.)

Default 7 — This session does not preempt any other session.

record

Syntax

[no] record

Context

config>router>mpls>lsp>primary

config>router>mpls>lsp>secondary

Description

This command enables recording of all the hops that an LSP path traverses. Enabling **record** increases the size of the PATH and RESV refresh messages for the LSP, since this information is carried end-to-end along the path of the LSP. The increase in control traffic per LSP may impact scalability.

The **no** form of this command disables the recording of all the hops for the given LSP. There are no restrictions as to when the **no** command can be used. 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

Description

This command enables recording of all the labels at each node that an LSP path traverses. Enabling the **record-label** command will also enable the **record** command, if it is not already enabled.

The **no** form of this command disables the recording of the hops that an LSP path traverses.

Default

record-label

srlg

Syntax

[no] srlg

Context

config>router>mpls>lsp>secondary

Description

This command enables the use of the SRLG constraint in the CSPF computation of a secondary path for an RSVP-TE LSP at the head-end LER. When this feature is enabled, CSPF includes the SRLG constraint in the computation of the secondary RSVP-TE LSP path if **path-computation-method local-cspf** is configured on the LSP.

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-TE task queries CSPF again, which provides the list of SRLG 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 path is set up. If CSPF does not find a path, MPLS/RSVP-TE keeps retrying the requests to CSPF.

If CSPF is not enabled on the LSP, 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 the primary path does not come up or is not configured, the SRLG secondary path 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, MPLS/RSVP-TE task signals it and the LSP uses it.

As soon as the primary path is configured and successfully established, MPLS/RSVP-TE moves the LSP to the primary path and signals all SRLG secondary paths.

Any time the primary path is reoptimized, has undergone a make-before-break (MBB) operation, or has come back up after being down, the MPLS/RSVP-TE task checks with CSPF to determine if the SRLG secondary path should be resignaled. If the MPLS/RSVP-TE 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 make-before-break 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 is, when the primary path goes down), the LSP uses of an eligible SRLG secondary path if the secondary path is in the up state. If all secondary eligible SRLG paths are in the down state, MPLS/RSVP-TE 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-TE 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-TE 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 its ineligibility status removed when any of the following events occurs:

- A successful MBB operation of the standby SRLG path occurs, making it eligible again.

- The standby path goes down, in which case MPLS/RSVP-TE 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 configured retry-limit value, 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.

Once the primary path of the LSP is set up and is operationally up, any subsequent changes to the SRLG 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 due to 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 reoptimization of the LSP path, or a change by the user to any of the path constraints.

Once an SRLG secondary path is set up and is operationally up, any subsequent changes to the SRLG 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 due to a failure, to a resigaling 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 a change by the user to any of the path constraints of the secondary path, including enabling or disabling the SRLG constraint itself.

In addition, any user-configured **include** or **exclude** admin group statements for this secondary path are checked along with the SRLG constraints by CSPF.

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

Default

no srlg

standby

Syntax

[no] standby

Context

config>router>mpls>lsp>secondary

Description

The secondary path LSP is normally signaled if the primary path LSP fails. The **standby** keyword ensures that the standby 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.



Note: A priority level can be assigned to standby secondary paths using the [path-preference](#) command.

The **no** form of this command specifies that the secondary LSP is signaled when the primary path LSP fails.

Default

n/a

3.24.2.1.9 LSP path commands

path

Syntax

[no] **path** *path-name*

Context

config>router>mpls

Description

This command creates 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 shut down, 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 the services that are actively using these LSPs will be affected. A path must be shut down 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.

Default

n/a

Parameters

path-name

specifies the unique case-sensitive alphanumeric name label for the LSP path, up to 32 characters in length

hop

Syntax

hop *hop-index* *ip-address*{**strict** | **loose**}

hop *hop-index***sid-label** *sid-value*

no hop *hop-index*

Context

```
config>router>mpls>path
```

Description

This command specifies 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, a loopback 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 addresses can be included as the first and last hop. A hop list can include the ingress interface IP address, the system IP address, and the egress IP address of any of the hops being specified.

When the **sid-label** parameter is specified, this command specifies an MPLS label value for a hop in the path of an SR-TE LSP. The label value implied by the SID is only used when the path is used by an SR-TE LSP. The **loose** and **strict** options do not apply for a hop containing a SID label; a SID hop must always be strict.

When the IP address is specified, the hop must be configured as either loose or strict. The **loose** option specifies that the route taken by the LSP from the previous hop to this hop can traverse other routers. The **strict** option specifies 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, then 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. For both options, multiple hop entries with the same IP address are flagged as errors.

The **no** form of this command deletes hop list entries for the path. All the LSPs currently using this path are affected. Additionally, all services actively using these LSPs are affected. The path must be first shut down in order to delete the hop from the hop list. The **no hop hop-index** command does not result in any action other than generating a warning message indicating that the path is administratively up.

Default

n/a

Parameters

hop-index

specifies the hop index, which is used to order the specified hops. 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, loopback, or network interface IP address of the transit router

sid-value

specifies any valid MPLS/SR label value. The value is not restricted by any locally defined label ranges.

Values 32 to 1048575

strict

specifies 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.

loose

specifies that the route taken by the LSP from the previous hop to this hop can traverse other routers

3.24.2.1.10 Static LSP commands

static-lsp

Syntax

[no] static-lsp *lsp-name*

Context

config>router>mpls

Description

This command configures static LSPs on the ingress router. The static LSP is a manually configured 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 before it can be deleted. 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.

Parameters

lsp-name

identifies the LSP. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

push

Syntax

push *label nexthop ip-address*

no push *label*

Context

config>router>mpls>static-lsp

Description

This command specifies the label to be pushed onto 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 31 are 7705 SAR reserved

Label values 32 through 1023 are available for static assignment

Label values 1024 through 2047 are reserved for future use

Label values 2048 through 18431 are statically assigned for services

Label values 28672 through 131071 are dynamically assigned for both MPLS and services

Label values 131072 through 1048575 are reserved for future use.

Values 16 to 1048575

ip-address

specifies the IP address of the next hop toward the LSP egress router. If an ARP entry for the next hop exists, then 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 send an ARP request for the configured next hop at fixed intervals.

to

Syntax

to *ip-address*

Context

config>router>mpls>static-lsp

Description

This command specifies the system IP address of the egress router for the static LSP. For LSPs that are used as transport tunnels for services, the **to ip-address** must be the system IP address. If the **to ip-address** does not match the SDP address, the LSP is not included in the SDP definition.

This command is required when creating an LSP.

Default

n/a

Parameters

ip-address

identifies the egress router system address

static-lsp-fast-retry

Syntax

static-lsp-fast-retry *seconds*

no static-lsp-fast-retry

Context

config>router>mpls

Description

This command specifies the fast-retry timer that can be configured for static LSPs. When a static LSP is trying to come up, MPLS tries to resolve the ARP entry for the next hop of the LSP. If the next hop is still down or unavailable, the request may fail. In that case, MPLS starts a non-configurable timer of 30 seconds before making the next request. The fast-retry timer allows the user to configure a shorter retry timer so that the LSP comes up shortly after the next hop is available.

Default

30

Parameters

seconds

fast-retry timer value, in seconds

Values 1 to 30

3.24.2.2 Configuration commands (RSVP-TE)

- [Generic commands](#)
- [RSVP-TE global commands](#)
- [Interface commands](#)
- [Message pacing commands](#)

3.24.2.2.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

```
config>router>rsvp  
config>router>rsvp>interface
```

Description

This command disables the RSVP-TE protocol instance or the RSVP-related functions for the interface. The RSVP-TE configuration information associated with this interface is retained. When RSVP-TE is administratively disabled, all the RSVP-TE sessions are torn down.

The **no** form of this command administratively enables RSVP-TE on the interface.

Default

shutdown

3.24.2.2.2 RSVP-TE global commands

```
rsvp
```

Syntax

[no] **rsvp**

Context

```
config>router
```

Description

This command creates the RSVP-TE protocol instance and enables RSVP-TE configuration.

RSVP-TE is enabled by default.

RSVP-TE is used to set up LSPs. RSVP-TE should be enabled on all router interfaces that participate in signaled LSPs.

The **no** form of this command deletes this RSVP-TE protocol instance and removes all configuration parameters for this RSVP-TE instance. To suspend the execution and maintain the existing configuration, use the **shutdown** command. RSVP-TE must be shut down before the RSVP-TE instance can be deleted. If RSVP-TE is not shut down, the **no rsvp** command does nothing except issue a warning message on the console indicating that RSVP-TE is still administratively enabled.

Default

no shutdown

```
entropy-label-capability
```

Syntax

[no] **entropy-label-capability**

Context

```
config>router>rsvp
```

Description

This command enables or disables the entropy label capability for an RSVP-TE LSP. When enabled, the egress LER (eLER) signals to the ingress LER (iLER) that the LSP is capable of using entropy labels.

The **no** form of the command disables entropy label capability.

Default

```
no entropy-label-capability
```

graceful-shutdown

Syntax

```
[no] graceful-shutdown
```

Context

```
config>router>rsvp
```

```
config>router>rsvp>interface
```

Description

This command initiates a graceful shutdown of the specified RSVP interface (referred to as a maintenance interface) or all RSVP interfaces on the node (referred to as a maintenance node). When this command is executed, the node performs the following operations in no specific order.

A PathErr message with an error sub-code of "Local Maintenance on TE Link required" is generated for each LSP that is in transit at this node and is using a maintenance interface as its outgoing interface. A PathErr message with the error code "Local node maintenance required" is generated if all interfaces are affected.

A single make-before-break attempt is performed for all adaptive CSPF LSPs that originate on the node and whose paths make use of the maintenance interfaces listed in the PathErr message. 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 resignals any bypass or detour LSP that uses the maintenance interfaces as soon as they are not active. The maintenance node floods an IGP TE LSA/LSP containing a Link TLV for the links under graceful shutdown with the Traffic Engineering metric set to 0xffffffff and the Unreserved Bandwidth parameter set to zero (0).

Upon receipt of the PathErr message, an intermediate LSR tears down and resignals any bypass LSP whose path makes use of the listed maintenance interfaces as soon as no associations with a protected LSP are active. The node does not take any action on a detour LSP whose path makes use of the listed maintenance interfaces.

Upon receipt of the PathErr message, a head-end LER 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, there are no alternative paths that can be found.
- an adaptive CSPF LSP whose path has explicit hops defined using the listed maintenance interfaces or node
- a CSPF LSP that has the adaptive option disabled and whose current path is over the listed maintenance interfaces in the PathErr message. These are not subject to make-before-break.
- a non-CSPF LSP whose current path is over the listed maintenance interfaces in the PathErr message

Upon receipt of the updated IPG TE LSA/LSP for the maintenance interfaces, the head-end LER updates the TE database. This information will be used at the next scheduled CSPF computation for any LSP whose path might traverse any of the maintenance interfaces.

The **no** form of the 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.

Default

n/a

implicit-null-label

Syntax

[no] implicit-null-label

Context

config>router>rsvp

Description

This command enables the implicit null label to be included in RESV messages sent by the egress LER (eLER) to the previous-hop LSR. The implicit null label is enabled for all LSPs for which the router is the eLER.

When the implicit null label is signaled to the LSR, it pops the outer label before sending the MPLS packet to the eLER; this is known as penultimate hop popping.

RSVP-TE must be shut down before this command can be used.

The **no** form of this command disables the signaling of the implicit null label.

Default

no implicit-null-label

keep-multiplier

Syntax

[no] keep-multiplier *number*

no keep-multiplier

Context

```
config>router>rsvp
```

Description

This command is used by RSVP-TE to declare that a reservation is down or the neighbor is down. The **keep-multiplier** *number* is used with the **refresh-time** command to determine when RSVP-TE will declare the session down.

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

Default

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
```

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.

Default

exclude

Parameters

include

the node-id sub-object is included in the RRO

exclude

the node-id sub-object is not included in the RRO

rapid-retransmit-time

Syntax

rapid-retransmit-time *hundred-milliseconds*

no rapid-retransmit-time

Context

config>router>rsvp

Description

This command is used to define the value of the rapid retransmission interval. This is used in the retransmission mechanism based on an exponential backoff timer in order to handle unacknowledged message-*_id* objects. The RSVP-TE message with the same message-id is retransmitted every 2 × rapid-retransmit-time interval. The node will stop retransmission of unacknowledged RSVP-TE 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 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

5 (which represents 500 msec)

Parameters

hundred-milliseconds

1 to 100, in units of 100 msec

rapid-retry-limit

Syntax

rapid-retry-limit *number*

no rapid-retry-limit

Context

config>router>rsvp

Description

This command is used to define the value of the rapid retry limit. This is used in the retransmission mechanism based on an exponential backoff timer in order to handle unacknowledged message-*_id* objects. The RSVP-TE message with the same message-*_id* is retransmitted every 2 × rapid-retransmit-time interval. The node will stop retransmission of unacknowledged RSVP-TE 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

3

Parameters

number

1 to 6, integer values

refresh-reduction-over-bypass

Syntax

refresh-reduction-over-bypass [**enable** | **disable**]

Context

config>router>rsvp

Description

This command enables the refresh reduction capabilities over all bypass tunnels originating on this 7705 SAR PLR node or terminating on this 7705 SAR Merge Point (MP) node.

By default, this is disabled. Since 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 that 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-TE messages pertaining to LSP paths tunneled over the bypass. It will also not send message-id in RSVP-TE messages. This effectively disables summary refresh.

Default

disable

refresh-time

Syntax

refresh-time *seconds*

no refresh-time

Context

config>router>rsvp

Description

This command controls the interval, in seconds, between the successive PATH and RESV refresh messages. RSVP-TE 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

30

Parameters

seconds

specifies the refresh time in seconds

Values 1 to 65535

3.24.2.2.3 Interface commands

interface

Syntax

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

Context

config>router>rsvp

Description

This command enables RSVP-TE protocol support on an IP interface. No RSVP-TE commands are executed on an IP interface where RSVP-TE is not enabled.

The **no** form of this command deletes all RSVP-TE commands such as **hello-interval** and **subscription**, which are defined for the interface. The RSVP-TE interface must be shut down 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.

Parameters

ip-int-name

specifies the network IP interface. The interface name cannot be in the form of an IP address. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

Values 1 to 32 alphanumeric characters

auth-keychain

Syntax

auth-keychain *name*
no auth-keychain

Context

config>router>rsvp>interface

Description

This command associates an authentication keychain with the RSVP-TE interface. The keychain is a collection of keys used to authenticate RSVP-TE messages from remote peers. The keychain allows the rollover of authentication keys during the lifetime of a session and also supports stronger authentication algorithms than clear text and MD5.

The keychain must already be defined in the **config>system>security>keychain** context.

Either the **authentication-key** command or the **auth-keychain** command can be used by RSVP-TE, but both cannot be supported at the same time. If both commands are configured, the **auth-keychain** configuration will be applied and the **authentication-key** command will be ignored.

By default, authentication is not enabled.

Default

no auth-keychain

Parameters

name

the name of an existing keychain, up to 32 characters

authentication-key

Syntax

authentication-key {*authentication-key*|*hash-key*}[*hash* | *hash2*]
no authentication-key

Context

config>router>rsvp>interface

Description

This command specifies the authentication key to be used between RSVP-TE neighbors to authenticate RSVP-TE messages. Authentication uses the MD5 message-based digest.

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

A 7705 SAR 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:

- the HMAC-MD5 authentication algorithm
- the key used with the authentication algorithm
- the lifetime of the key; the user-entered key is valid until the user deletes it from the interface
- the source address of the sending system
- the latest sending sequence number used with this key identifier

A 7705 SAR RSVP-TE sender transmits an authenticating digest of the RSVP-TE message, computed using the shared authentication key and a keyed hash algorithm. The message digest is included in an integrity object that also contains a flags field, a key identifier field, and a sequence number field. The 7705 SAR RSVP-TE sender complies with the procedures for RSVP-TE message generation in RFC 2747, *RSVP Cryptographic Authentication*.

A 7705 SAR RSVP-TE receiver uses the key together with the authentication algorithm to process received RSVP-TE 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-TE messages sent over the bypass tunnel. If the PLR receives an RSVP-TE message with an integrity object, it will perform the digest verification for the key of the interface over which the packet was received. If this fails, the packet is dropped. If the received RSVP-TE message is an RESV message and does not have an integrity object, then the PLR node will accept it only if it originated from the MP node.

A 7705 SAR MP node will accept RSVP-TE messages received over the bypass tunnel with and without the integrity object. If an integrity object is present, the proper digest verification for the key of the interface over which the packet was received is performed. If this fails, the packet is dropped.

The 7705 SAR MD5 implementation does not support the authentication challenge procedures in RFC 2747.

Either the **authentication-key** command or the **auth-keychain** command can be used by RSVP-TE, but both cannot be supported at the same time. If both commands are configured, the **auth-keychain** configuration will be applied and the **authentication-key** command will be ignored.

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

Default

no authentication-key – the authentication key value is the null string

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 (such as #, \$, or spaces), 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

specifies the key is entered 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

specifies the key is entered in a more complex encrypted form. If the **hash2** keyword is not used, the less-encrypted **hash** form is assumed.

bfd-enable**Syntax**

[no] bfd-enable

Context

config>router>rsvp>interface

Description

This command enables the use of bidirectional forwarding (BFD) to control the state of the associated RSVP-TE interface. This causes RSVP-TE 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>if>bfd** context.

The BFD session on the interface may already have been started because of a prior registration with another protocol; for example, OSPF or IS-IS.

The registration of an RSVP-TE 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 due to not receiving an 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-TE with BFD is removed as soon as the last RSVP-TE session is cleared.

The registration of an RSVP-TE interface with BFD is performed independently of whether RSVP-TE hello is enabled on the interface or not. However, hello timeout clears all sessions toward the neighbor and RSVP-TE deregisters with BFD at the clearing of the last session.

An RSVP-TE session is associated with a neighbor based on the interface address that the PATH message is sent to. If multiple interfaces exist to the same node, each interface is treated as a separate RSVP-TE neighbor. The user must enable BFD on each interface, and RSVP-TE 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-TE 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-TE protocol adjacency.

Default

no bfd-enable

hello-interval

Syntax

hello-interval *milli-seconds*

no hello-interval

Context

config>router>rsvp>interface

Description

This command configures the time interval between RSVP-TE hello messages.

RSVP-TE hello packets are used to detect loss of RSVP-TE connectivity with the neighboring node. Hello packets detect the loss of a neighbor more quickly than it would take for the RSVP-TE session to time out based on the refresh interval. After the loss of the of **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

3000

Parameters

milli-seconds

specifies the RSVP-TE hello interval in milliseconds, in multiples of 1000. A 0 (zero) value disables the sending of RSVP-TE 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>interface

Description

This command enables or disables the use of the implicit null label over a specific RSVP-TE interface.

The implicit null label is enabled for all LSPs for which the router is the eLER and for which the PATH message is received from the previous-hop LSR over the RSVP-TE interface.

With facility backup, if the eLER is also the merge point (MP) node, the incoming interface for the PATH refresh message over the bypass tunnel dictates whether the packet will use the implicit null label. Similarly, with one-to-one backup, if the eLER is also the detour merge point (DMP) node, the incoming interface for the PATH refresh message over the detour LSP dictates whether the packet will use the implicit null label.

By default, an RSVP-TE interface inherits the RSVP-TE level configuration.

The interface must be shut down before this command can be used.

The **no** form of this command resets the interface to the RSVP-TE level configuration.

Default

no implicit-null-label

Parameters

enable

enables the implicit null label on the interface

disable

disables the implicit null label on the interface

refresh-reduction

Syntax

[no] refresh-reduction

Context

config>router>rsvp>interface

Description

This command enables the use of the RSVP-TE overhead refresh reduction capabilities on this RSVP-TE interface.

When this option is enabled, a 7705 SAR node will enable support for three capabilities:

- it will accept bundle RSVP-TE messages from its peer over this interface
- it will attempt to perform reliable RSVP-TE message delivery to its peer
- it will use summary refresh messages to refresh PATH and RESV states

The reliable message delivery must be explicitly enabled by the user after refresh reduction is enabled. The other two capabilities are enabled immediately.

A bundle RSVP-TE message is intended to reduce the overall message handling load. A bundle message consists of a bundle header followed by one or more bundle sub-messages. A sub-message can be any regular RSVP-TE message except another bundle message. A 7705 SAR node will only process received bundle RSVP-TE messages but will not generate them.

When reliable RSVP-TE message delivery is supported by both the node and its peer over the RSVP-TE interface, an RSVP-TE message is sent with a `message_id` object. A `message_id` object can be added to any RSVP-TE message when sent individually or as a sub-message of a bundle message.

If the sender sets the `ack_desired` flag in the `message_id` object, the receiver acknowledges the receipt of the RSVP-TE message by piggy-backing a `message_ack` object to the next RSVP-TE message it sends to its peer. Alternatively, an ACK message can also be used to send the `message_ack` object. In both cases, one or many `message_ack` objects could be included in the same message.

The 7705 SAR supports the sending of separate ACK messages only, but is capable of processing received `message_ack` objects piggy-backed to hop-by-hop RSVP-TE messages, such as PATH and RESV.

The 7705 SAR sets the `ack_desired` flag only in non-refresh RSVP-TE messages and in refresh messages that contain new state information.

A retransmission mechanism based on an exponential backoff timer is supported in order to handle unacknowledged `message_id` objects. The RSVP-TE message with the same `message_id` is retransmitted every $2 \times$ 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.

There is also a maximum number of retransmissions of an unacknowledged RSVP-TE message **`rapid-retry-limit`**. The node will stop retransmission of unacknowledged RSVP-TE messages whenever the updated backoff interval exceeds the value of the regular **`refresh-time`** interval or the number of retransmissions reaches the value of the **`rapid-retry-limit`** parameter, whichever comes first. These two parameters are configurable globally on a system in the **`config>router>rsvp`** context.

Summary refresh consists of sending a summary refresh message containing a `message_id` list object. The fields of this object are populated each with the value of the `message_identifier` field in the `message_id` object of a previously sent individual PATH or RESV message. The summary refresh message is sent every refresh regular interval as 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 as if a regular PATH or RESV refresh message was received from the peer. If a specific `message_identifier` field does not match, then the node sends a `message_id_nack` object to the originator of the message.

The above capabilities are referred to collectively as "refresh overhead reduction extensions". When the refresh-reduction is enabled on a 7705 SAR RSVP-TE interface, the node indicates this to its peer by setting a "refresh-reduction-capable" bit in the flags field of the common RSVP-TE header. If both peers of an RSVP-TE interface set this bit, all the above three capabilities can be used. Furthermore, the node monitors the settings of this bit in received RSVP-TE messages from the peer on the interface. As soon as this bit is cleared, the 7705 SAR stops sending summary refresh messages. If a peer did not set the "refresh-reduction-capable" bit, a node does not attempt to send summary refresh messages.

However, if the peer did not set the "refresh-reduction-capable" bit, then a node with refresh reduction enabled and reliable message delivery enabled will still attempt to perform reliable message delivery with this peer. If the peer does not support the `message_id` object, it returns the error message "unknown object class". In this case, the 7705 SAR node retransmits the RSVP-TE message without the `message_id` object and reverts to using this method for future messages destined for this peer.

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

Default

`no refresh-reduction`

reliable-delivery

Syntax

[no] **reliable-delivery**

Context

config>router>rsvp>if>refresh-reduction

Description

This command enables reliable delivery of RSVP-TE messages over the RSVP-TE interface. When **refresh-reduction** is enabled on an interface and **reliable-delivery** is disabled, then the 7705 SAR will send a message_id and not set ACK desired in the RSVP-TE messages over the interface.

The 7705 SAR does not expect an ACK but will accept it if received. The node will also accept message ID and reply with an ACK when requested. In this case, if the neighbor set the "refresh-reduction-capable" bit in the flags field of the common RSVP-TE header, the node will enter 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-TE message pacing is disabled on all RSVP-TE 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 CLI. When the **msg-pacing** option is enabled, the user cannot enable the **reliable-delivery** option on any interface on this system. An error message will also be generated in CLI after such an attempt.

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

Default

no reliable-delivery

subscription

Syntax

subscription *percentage*

no subscription

Context

config>router>rsvp>interface

Description

This command configures the percentage of the link bandwidth that RSVP-TE can use for reservation and sets a limit for the amount of over-subscription or under-subscription allowed on the interface.

When the **subscription** is set to zero, no new sessions are permitted on this interface. If the percentage is exceeded, the reservation is rejected and a log message is generated.

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

Default

100

Parameters*percentage*

specifies the percentage of the interface's bandwidth that RSVP-TE allows to be used for reservations

Values 0 to 1000**3.24.2.2.4 Message pacing commands****msg-pacing****Syntax****[no] msg-pacing****Context**

config>router>rsvp

Description

This command enables RSVP-TE message pacing, which is defined by the **max-burst** and **period** commands. 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

Description

This command specifies the maximum number of RSVP-TE messages that can be sent under normal operating conditions, as specified by the **period** command. The **no** form of this command reverts to the default value.

Default

650

Parameters*number*

maximum number of RSVP-TE messages

Values 100 to 1000, in increments of 10

period

Syntax**period** *milli-seconds***no period****Context**

config>router>rsvp>msg-pacing

Description

This command specifies the time interval, in milliseconds, during which the router can send RSVP-TE messages, as specified by the **max-burst** command. The **no** form of this command reverts to the default value.

Default

100

Parameters*milli-seconds*

the time interval during which the router can send RSVP-TE messages

Values 100 to 1000 milliseconds, in increments of 10 milliseconds

3.24.2.3 Show commands (MPLS)



Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

admin-group

Syntax**admin-group** *group-name*

Context

show>router>mpls

Description

This command displays MPLS administrative group information.

Parameters

group-name
specifies the administrative group name

Output

The following output is an example of MPLS administrative group information, and [Table 9: Router MPLS admin-group field descriptions](#) describes the fields.

Output example

```
A:ALU-1# show router mpls admin-group
=====
MPLS Administrative Groups
=====
Group Name                Group Value
-----
green                     15
red                       25
yellow                    20
-----
No. of Groups: 3
=====
A:ALU-1#
```

Table 9: Router MPLS admin-group field descriptions

Label	Description
Group Name	The name of the administrative group. The name identifies the administrative group within a router instance.
Group Value	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	The total number of configured administrative groups within the router instance

bypass-tunnel

Syntax

bypass-tunnel [*to ip-address*] [**protected-lsp** [*lsp-name*]] [**dynamic** | **manual**] [**detail**]

Context

show>router>mpls

Description

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 that serves as a backup for a set of LSPs. This type of 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**
displays dynamically assigned labels for bypass protection
- manual**
displays manually assigned labels for bypass protection
- detail**
displays detailed information

Output

The following output is an example of MPLS bypass tunnel information, and [Table 10: Router MPLS bypass-tunnel field descriptions](#) describes the fields.

Output example

```
A:ALU-12>show>router>mpls# bypass-tunnel to 10.20.1.4
=====
Legend :  m - Manual                d - Dynamic
=====
To          State    Out I/F  Out Label  Reserved   Protected   Type
              BW (Kbps)   LSP Count
-----
10.20.1.4    Up        lag      *-*        131071      0
-----
Bypass Tunnels : 1
=====
A:ALU-12>show>router>mpls#
```

Table 10: Router MPLS bypass-tunnel field descriptions

Label	Description
To	The system IP address of the egress router
State	The administrative state of the LSP
Out I/F	The name of the network IP interface

Label	Description
Out Label	The incoming MPLS label on which to match
Reserved BW (Kbps)	The amount of bandwidth in kilobytes per second reserved for the LSP
Protected LSP Count	The number of times this LSP has used a protected LSP
Type	The type of protected LSP

interface

Syntax

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

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

Context

show>router>mpls

Description

This command displays MPLS interface information.

Parameters

ip-int-name

identifies the network IP interface. The interface name cannot be in the form of an IP address. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

ip-address

specifies the system or network interface IP address

label-map label

specifies the MPLS label on which to match

Values 32 to 1023

statistics

displays IP address and the number of packets and octets sent and received on an interface basis

Output

The following output is an example of MPLS interface information, and [Table 11: Router MPLS interface field descriptions](#) describes the fields.

Output example

```
ALU-12# show router mpls interface
```

```
=====
```

MPLS Interfaces

```
=====
```

Interface	Port-id	Adm	Opr	TE-metric
system	vport-1	Up	Up	None
Admin Groups	None			
Srlg Groups	None			
ip-10.10.1.2	1/1/1	Up	Up	None
Admin Groups	None			
Srlg Groups	None			
ip-10.10.4.2	1/1/2	Up	Up	None
Admin Groups	None			
Srlg Groups	None			
ip-10.10.3.2	1/1/3	Up	Up	None
Admin Groups	None			
Srlg Groups	None			

```
-----
```

Interfaces : 4

```
=====
```

```
*A:ALU-48>config>router>mpls# show router mpls interface "to-104" label-map 35
```

```
=====
```

MPLS Interface : to-104 (Label-Map 35)

```
=====
```

In Label	In I/F	Out Label	Out I/F	Next Hop	Type	Adm	Opr
35	1/1/1	n/a	n/a	n/a	Static	Up	Down

```
-----
```

Interfaces : 1

```
=====
```

```
*A:ALU-48>config>router>mpls#
```

```
ALU-12# show router mpls interface statistics
```

```
=====
```

MPLS Interface (statistics)

```
=====
```

Interface	: ip-10.10.1.1	
Transmitted	: Pkts - 6	Octets - 540
Received	: Pkts - 0	Octets - 0
Invalid	: Labels - 0	
Invalid	: IPoMPLS Pkts - 0	
Invalid	: Stack Too Big Pkts - 0	
Invalid	: TTL Expired Pkts - 0	
Invalid	: Other Discard Pkts - 0	
Last Invalid	: Label Value - 0	
Last Invalid	: Label Position - 0	
Interface	: ip-10.10.2.1	
Transmitted	: Pkts - 0	Octets - 0
Received	: Pkts - 0	Octets - 0
Invalid	: Labels - 0	
Invalid	: IPoMPLS Pkts - 0	
Invalid	: Stack Too Big Pkts - 0	
Invalid	: TTL Expired Pkts - 0	
Invalid	: Other Discard Pkts - 0	
Last Invalid	: Label Value - 0	
Last Invalid	: Label Position - 0	

```
=====
```

ALU-12#

Table 11: Router MPLS interface field descriptions

Label	Description
Interface	The interface name
Port-id	The port ID in the <i>slot/mda/port</i> format
Adm	The administrative state of the interface
Opr	The operational state of the interface
Te-metric	The traffic engineering metric used on the interface
Srlg Groups	The shared risk link group (SRLG)
Interfaces	The total number of interfaces
Transmitted	The number of packets and octets transmitted from the interface
Received	The number of packets and octets received
In Label	The ingress label
In I/F	The ingress interface
Out Label	The egress label
Out I/F	The egress interface
Next Hop	The next-hop IP address for the static LSP
Type	Indicates whether the label value is statically or dynamically assigned
Invalid	Labels – the number of incoming packets discarded due to invalid labels
	IPoMPLS Pkts – the number of incoming labeled packets discarded due to invalid IP packet headers in the packet
	Stack Too Big Pkts – the number of incoming packets discarded due to having greater than the maximum number of labels in the label stack (that is, greater than five)
	TTL Expired Pkts – the number of incoming packets discarded due to exceeding the maximum Time-To-Live (TTL) value
	Other Discard Pkts – the number of incoming packets discarded due to internal errors (for example, memory corruption or invalid label table programming)
Last Invalid	Label Value – the value of the last invalid label received

Label	Description
	Label Position – the position in the label stack of the last invalid label received

lsp

Syntax

lsp [*lsp-name*] [**status** {**up** | **down**}] [**from** *ip-address* | **to** *ip-address*] [**detail**] [**auto-lsp** {**all** | **mesh-p2p** | **one-hop-p2p**}]

lsp {**transit** | **terminate**} [**status** {**up** | **down**}] [**from** *ip-address* | **to** *ip-address* | *lsp-name*] [**detail**]

lsp count

lsp [*lsp-name*] **activepath** [**auto-lsp** {**all** | **mesh-p2p** | **one-hop-p2p**}]

lsp [*lsp-name*] **path** [*path-name*] [**status** {**up** | **down**}] [**detail**] [**auto-lsp** {**all** | **mesh-p2p** | **one-hop-p2p**}]

lsp [*lsp-name*] **path** [*path-name*] **mbb** [**auto-lsp** {**all** | **mesh-p2p** | **one-hop-p2p**}]

Context

show>router>mpls

Description

This command displays LSP details.

Parameters

lsp-name

specifies the name of the LSP used in the path

status up

displays an LSP that is operationally up

status down

displays 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 LSPs that transit the router

terminate

displays the LSPs that terminate at the router

name

displays the IP address of the named LSP

count

displays the total number of LSPs

activepath

displays the present path being used to forward traffic

path-name

specifies the name of the path carrying the LSP

mbb

displays make-before-break (MBB) information

detail

displays detailed information

auto-lsp {all | mesh-p2p | one-hop-p2p}

specifies the type of auto-LSP or all auto-LSPs

Output

The following outputs are examples of MPLS LSP information:

- MPLS LSP ([Output example](#), [Table 12: Router MPLS LSP field descriptions](#))
- MPLS LSP Detail ([Output example](#), [Table 13: MPLS LSP detail field descriptions](#))
- MPLS LSP Path Detail ([Output example](#), [Table 14: Router MPLS LSP path detail field descriptions](#))
- MPLS LSP Path MBB ([Output example](#), [Table 15: Router MPLS LSP path MBB field descriptions](#))
- MPLS Auto-LSP ([Output example](#), [Table 16: Router MPLS auto-LSP field descriptions](#))

Output example

```
A:ALU-48# show router mpls lsp
```

```
=====
```

```
MPLS LSPs (Originating)
```

```
=====
```

LSP Name	To	Fastfail Config	Adm	Opr
to-104	10.10.10.104	Yes	Up	Up
to-103	10.0.0.0	Yes	Up	Up
to-99	10.10.10.99	No	Up	Up
to-100	10.10.10.100	No	Up	Up
to-49	10.20.30.49	No	Dwn	Up

```
-----
```

```
LSPs : 5
```

```
=====
```

```
A:ALU-48#
```

```
*A:ALU-48# show router mpls lsp to-104
```

```
=====
```

```
MPLS LSPs (Originating)
```

```
=====
```

LSP Name	To	Fastfail Config	Adm	Opr
to-104	10.10.10.104	Yes	Up	Dwn

```
-----
```

```
LSPs : 1
```

```
=====
```

*A:ALU-48#

Table 12: Router MPLS LSP field descriptions

Label	Description
LSP Name	The name of the LSP used in the path
To	The system IP address of the egress router for the LSP
FastFail Config	enabled – fast reroute is enabled. In the event of a failure, traffic is immediately rerouted on the precomputed protection LSP to minimize packet loss
	disabled – there is no protection LSP from each node on the primary path
Adm State	Down – the path is administratively disabled
	Up – the path is administratively enabled
Oper State	Down – the path is operationally down
	Up – the path is operationally up
LSPs	The total number of LSPs configured

Output example

```
# show router mpls lsp detail
=====
MPLS LSPs (Originating) (Detail)
=====
Legend :
+ - Inherited
=====
Type : Originating
-----
LSP Name      : C_F_1
LSP Type      : RegularLsp          LSP Tunnel ID      : 1
LSP Index     : 1                  TTM Tunnel Id      : 1
From          : 10.20.1.3
To            : 10.20.1.6
Adm State     : Up
LSP Up Time   : 0d 00:07:27
Transitions   : 3
Retry Limit   : 0
Signaling     : RSVP
Hop Limit     : 255
Adaptive      : Enabled
FastReroute   : Disabled
PathCompMethod : pce
FallbkPathComp : none
Metric        : N/A
Load Bal Wt   : N/A
Include Grps  :
None
Least Fill    : Disabled

Oper State    : Up
LSP Down Time : 0d 00:00:00
Path Changes  : 3
Retry Timer   : 20 sec
Resv. Style   : SE
Negotiated MTU : 1500
ClassType     : 0
Oper FR       : Disabled
ADSPEC        : Enabled

Metric Type   : igp
ClassForwarding : Disabled
Exclude Grps  :
None
```

```

BFD Template      : None           BFD Ping Intvl      : 60
BFD Enable        : False          BFD Failure-action  : None
WaitForUpTimer    : 4

Revert Timer      : Disabled        Next Revert In      : N/A
Entropy Label     : Enabled+        Oper Entropy Label  : Enabled
Negotiated EL     : Disabled
Auto BW           : Disabled
LdpOverRsvp       : Enabled
VprnAutoBind      : Enabled
IGP Shortcut      : Enabled         BGP Shortcut        : Enabled
IGP LFA           : Disabled        IGP Rel Metric      : Disabled
BGPTransTun       : Enabled
Oper Metric       : 100
Prop Adm Grp      : Disabled
PCE Report        : Enabled
PCE Control       : Enabled
Path Profile      : None
Admin Tags        : None
Lsp Self Ping     : Disabled+       Self Ping Timeouts  : 0
SelfPingOAMFail* : 0

Primary(a)        : C_F_1           Up Time             : 0d 00:07:28
Bandwidth         : 0 Mbps
=====
* indicates that the corresponding row element may have been truncated.

```

Table 13: MPLS LSP detail field descriptions

Label	Description
LSP Name	The name of the LSP used in the path
LSP Type	The type of LSP used in the path
LSP Tunnel ID	The LSP tunnel identifier
LSP Index	The LSP index value
TTM Tunnel Id	The tunnel table manager tunnel identifier
From	The IP address of the ingress router for the LSP
To	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
Oper State	Down – the path is operationally down
	Up – the path is operationally up
LSP Up Time	The length of time that the LSP has been operational
LSP Down Time	The total time that the LSP has not been operational
Transitions	The number of transitions that have occurred for the LSP

Label	Description
Path Changes	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 Limit	The number of attempts that the software should make to re-establish the LSP after it has failed
Retry Timer	The time, in seconds, for LSP re-establishment attempts after an LSP failure
Signaling	The signaling style
Style	SE (shared explicit) – 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 (fixed filter) – specifies a dedicated reservation environment with an explicit reservation scope. This reservation style specifies an explicit list of senders and a distinct reservation for each of them.
Hop Limit	The maximum number of hops that an LSP can traverse, including the ingress and egress routers
MTU	The size of the maximum transmission unit (MTU) that is negotiated during establishment of the LSP
Adaptive	Indicates whether the make-before-break option is enabled or disabled for resigaled paths
ClassType	The class type value
FastReroute	Enabled – fast reroute is enabled. In the event of a failure, traffic is immediately rerouted on the precomputed protection LSP, to minimize packet loss.
	Disabled – there is no protection LSP from each node on the primary path
Oper FR	Indicates whether FRR has been enabled or disabled
PathCompMethod	pce – PCE path computation method is configured local-cspf – local CSPF path computation method is configured
FallbkPathComp	none – no fallback method is configured local-cspf – local CSPF fallback is configured
CSPF	Indicates whether CSPF has been enabled or disabled
Metric	The TE metric value

Label	Description
Load Bal Wt	The load balancing weight value
Include Grps	The admin groups that are to be included by an LSP when signaling a path
Exclude Grps	The admin groups that are to be avoided by an LSP when signaling a path
Least Fill	Indicates whether the least fill option is enabled or disabled
Entropy Label	Indicates whether the entropy label capability is enabled or disabled
Oper Entropy Label	Indicates whether the entropy label capability is operational or not operational
Negotiated EL	Indicates whether the negotiated entropy label capability is enabled or disabled
Auto BW	Indicates whether auto-bandwidth adjustment is enabled or disabled
VprnAutoBind	Indicates whether VPRN autobind tunneling is enabled or disabled
IGP Shortcut	Indicates whether the IGP shortcut option is enabled or disabled
IGP LFA	Indicates whether the IGP loop-free alternate option is enabled or disabled
IGP Rel Metric	Indicates whether the IGP relative metric option is enabled or disabled
BGPTransTun	Indicates whether the BGP transport tunnel option is enabled or disabled
Oper Metric	The LSP operational metric value
Prop Adm Grp	Indicates whether the propagate administrative group option is enabled or disabled
PCE Report	Specifies the PCE reporting mode, either enabled, disabled, or inherited
PCE Control	Indicates whether the PCE control option is enabled or disabled
Primary (a)	The preferred path for the LSP
Up Time	The length of time that the path has been up
Bandwidth	The amount of bandwidth in megabits per second (Mbps) reserved for the LSP path

Output example

```
# show router mpls lsp path detail
=====
MPLS LSP Path (Detail)
=====
Legend :
  @ - Detour Available          # - Detour In Use
  b - Bandwidth Protected       n - Node Protected
  s - Soft Preemption
  S - Strict                    L - Loose
  A - ABR                      + - Inherited
=====
-----
LSP C_F_1
Path C_F_1
-----
LSP Name      : C_F_1
From          : 10.20.1.3
To            : 10.20.1.6
Admin State   : Up              Oper State      : Up
Path Name     : C_F_1
Path LSP ID   : 49666           Path Type      : Primary
Path Admin    : Up              Path Oper      : Up
Out Interface : 1/1/1           Out Label     : 524248
Path Up Time  : 0d 00:07:18     Path Down Time : 0d 00:00:00
Retry Limit   : 0               Retry Timer    : 20 sec
Retry Attempt : 0               Next Retry In  : 0 sec

BFD Configuration and State
Template      : None             Ping Interval   : 60
Enable       : False            State           : notApplicable
WaitForUpTimer : 4 sec          OperWaitForUpTimer: N/A
WaitForUpTmLeft : 0 sec

Adspec       : Enabled           Oper Adspec     : Enabled
PathCompMethod : pce             OperPathCompMethod: pce
MetricType    : igp              Oper MetricType  : igp
Least Fill    : Disabled         Oper LeastFill   : Disabled
FRR           : Disabled         Oper FRR         : Disabled
Propagate Adm Grp: Disabled      Oper Prop Adm Grp: Disabled
Inter-area    : N/A

PCE Report    : Enabled           Oper PCE Report  : Enabled
PCE Control   : Enabled           Oper PCE Control : Enabled
PCE Update ID : 0

Neg MTU       : 1500             Oper MTU         : 1500
Bandwidth     : No Reservation    Oper Bandwidth   : 0 Mbps
Hop Limit     : 255              Oper HopLimit    : 255
Record Route  : Record           Oper Record Route: Record
Record Label  : Record           Oper Record Label: Record
Setup Priority : 7                Oper SetupPriority: 7
Hold Priority  : 0                Oper HoldPriority : 0
Class Type    : 0                Oper CT          : 0
Backup CT     : None
MainCT Retry  : n/a
  Rem         :
MainCT Retry  : 0
  Limit       :
Include Groups :
None
Exclude Groups :
None
```

```

Adaptive       : Enabled           Oper Metric      : 100
Preference    : n/a
Path Trans     : 3                 CSPF Queries   : 0
Failure Code   : noError
Failure Node   : n/a
Explicit Hops  :
  No Hops Specified
Actual Hops    :
  10.10.2.3(10.20.1.3)             Record Label   : N/A
-> 10.10.2.1(10.20.1.1)           Record Label   : 524248
-> 10.10.1.2(10.20.1.2)           Record Label   : 524246
-> 10.10.4.4(10.20.1.4)           Record Label   : 524246
-> 10.10.9.6(10.20.1.6)           Record Label   : 524248
Computed Hops  :
  10.10.2.1(S)
-> 10.10.1.2(S)
-> 10.10.4.4(S)
-> 10.10.9.6(S)
Resignal Eligible: False
Last Resignal  : n/a              CSPF Metric     : 100
=====

```

Table 14: Router MPLS LSP path detail field descriptions

Label	Description
LSP Name	The name of the LSP used in the path
Path LSP ID	The LSP ID for the path
From	The IP address of the ingress router for the LSP
To	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
Oper State	Down – the path is operationally down
	Up – the path is operationally up
Path Name	The alphanumeric name of the path
Path Type	The type of path: primary or secondary
Path Admin	The administrative status of the path
Path Oper	The operational status of the path
OutInterface	The output interface of the LSP
Out Label	The output label of the LSP
Path Up Time	The length of time that the path has been operationally up
Path Down Time	The length of time that the path has been operationally down

Label	Description
Retry Limit	The number of times an LSP will retry before giving up
Retry Timer	The length of time between LSP signaling attempts
Retry Attempt	The number of attempts that have been made to re-establish the LSP
Next Retry	The time when the next attempt to re-establish the LSP will occur
PathCompMethod	pce – PCE path computation method is configured local-cspf – local CSPF path computation method is configured
OperPathCompMethod	Displays the path computation method in use (pce, local, or none)
Bandwidth	The amount of bandwidth in megabits per second (Mbps) reserved for the LSP
Oper Bandwidth	The bandwidth reserved by the LSP
Hop Limit	The limit on the number of hops taken by the LSP
Record Route	Indicates whether a list of routers for the LSP has been recorded
Record Label	Indicates whether a list of router labels has been recorded
Oper MTU	The operational MTU of the connection to the next hop
Neg MTU	The MTU negotiated between the router and its next hop
Adaptive	Indicates whether make-before-break is enabled or disabled for resigned paths
Include Grps	The admin groups that are to be included by an LSP when signaling a path
Exclude Grps	The admin groups that are to be avoided by an LSP when signaling a path
Preference	The path-preference priority number assigned to a standby secondary LSP path
Path Trans	The number of times a path has made a transition between up and down states
CSPF Queries	The number of requests made by the LSP to the TE database
Failure Code	The reason code for in-progress MBB failure. A value of none indicates that no failure has occurred.

Label	Description
Failure Node	The IP address of the node in the LSP at which the in-progress MBB failed. If no failure has occurred, this value is none .
Explicit Hops	The hops that have been specified by the user
Actual Hops	The hops that the route has taken, either numbered or unnumbered
Record Label	The label recorded at the specified hop
Computed Hops	The hops computed and returned from the routing database, either numbered or unnumbered
LastResignalAttempt	The system up time when the last attempt to resignal this LSP was made
Last Resignal	The last time the route was resigaled
Metric	The value of the metric
Last MBB	The header for the last make-before-break (MBB) information
MBB Type	An enumerated integer that specifies the type of make-before-break (MBB) operation. If none displays, then there is no MBB in progress or no last MBB.
MBB State	The state of the most recent invocation of the make-before-break functionality
Ended at	The system up time when the last MBB ended
Old Metric	The cost of the traffic engineered path for the LSP prior to MBB
In Progress MBB	Header for the currently in-progress MBB information
MBB Type	An enumerated integer that specifies the type of make-before-break (MBB) operation. If none displays, then there is no MBB in progress or no last MBB.
NextRetryIn	The amount of time remaining, in seconds, before the next attempt is made to retry the in-progress MBB
Started At	The time the current MBB began
RetryAttempt	The number of attempts for the MBB in progress
Failure Code	The reason code for in-progress MBB failure. A value of none indicates that no failure has occurred.

Label	Description
Failure Node	The IP address of the node in the LSP at which the in-progress MBB failed. If no failure has occurred, this value is none .

Output example

```
*A:ALU-48# show router mpls lsp path mbb
=====
MPLS LSP Path MBB
=====
LSP 1 Path 1
-----
LastResignalAttempt: 2008/04/08 11:42:33.22 PST   CSPF Metric   : 0

Last MBB:
MBB Type      : Timer-based Resignal                MBB State      : Success/Failed
Ended at      : 2008/04/08 11:12:23.76 PST          Old Metric     : 3000

In Progress MBB:
MBB Type      : Config Change                      NextRetryIn    : 16 sec
Started at    : 2008/04/08 12:01:02.20 PST          RetryAttempt   : 3
Failure Code: noCspfRouteToDestination             Failure Node   : 10.20.1.1

-----
LSP 2 Path 1
-----
LastResignalAttempt: 2008/04/08 11:42:33.54 PST   CSPF Metric   : 0

Last MBB:
MBB Type      : Timer-based Resignal                MBB State      : Success/Failed
Ended at      : 2008/04/08 11:12:24.76 PST          Old Metric     : 2000

-----
LSP 4 Path 1
-----
LastResignalAttempt: 2008/04/08 11:42:34.12 PST   CSPF Metric   : 0

In Progress MBB:
MBB Type      : Global Revertive                   NextRetryIn    : 10 sec
Started at    : 2008/04/08 11:45:02.20 PST          RetryAttempt   : 2
Failure Code: noCspfRouteToDestination             Failure Node   : 10.20.1.1
=====
*A:ALU-48#
```

Table 15: Router MPLS LSP path MBB field descriptions

Label	Description
LastResignalAttempt	The system up time when the last attempt to resignal this LSP was made
CSPF Metric	The value of the CSPF metric
Last MBB	Header for the last make-before-break (MBB) information

Label	Description
MBB Type	An enumerated integer that specifies the type of make-before-break (MBB) operation. If none displays, then there is no MBB in progress or no last MBB.
MBB State	The state of the most recent invocation of the make-before-break functionality
Ended at	The system up time when the last MBB ended
Old Metric	The cost of the traffic-engineered path for the LSP path prior to MBB
In Progress MBB	The header for the currently in-progress MBB information
MBB Type	An enumerated integer that specifies the type of make-before-break (MBB) operation. If none displays, then there is no MBB in progress or no last MBB.
NextRetryIn	The amount of time remaining, in seconds, before the next attempt is made to retry the in-progress MBB
Started At	The time that the current MBB began
RetryAttempt	The number of attempts for the MBB in progress
Failure Code	The reason code for in-progress MBB failure. A value of none indicates that no failure has occurred.
Failure Node	The IP address of the node in the LSP path at which the in-progress MBB failed. When no failure has occurred, this value is none .

Output example

```

A:ALU-48# show router mpls lsp auto-lsp mesh-p2p
=====
MPLS LSPs (Originating)
=====
LSP Name                               Type                Fastfail   Admin   Oper
Config                                State      State
-----
MESH-192.0.2.8-61456                   MeshP2P            Yes        Up      Up
MESH-192.0.2.9-61457                   MeshP2P            Yes        Up      Up
-----
Auto-LSPs : 2
=====
A:ALU-48#

```

Table 16: Router MPLS auto-LSP field descriptions

Label	Description
LSP Name	The name of the LSP used in the path

Label	Description
Type	The type of auto-LSP
FastFail Config	enabled – fast reroute is enabled. In the event of a failure, traffic is immediately rerouted on the precomputed protection LSP, to minimize packet loss
	disabled – there is no protection LSP from each node on the primary path
Admin State	Down – the path is administratively disabled
	Up – the path is administratively enabled
Oper State	Down – the path is operationally down
	Up – the path is operationally up
LSPs	The total number of LSPs configured

lsp-egress-stats

Syntax

lsp-egress-stats [*type* *lsp-type*] [*active*] [*template-match*]

lsp-egress-stats **lsp** *lsp-name*

Context

show>router>mpls

Description

This command displays RSVP LSP egress statistical information.

Parameters

lsp-type

specifies the type of LSP to display. The only available options are **p2p** and **p2mp**.

active

displays information from all LSPs with statistics collection enabled

template-match

displays information for a one-hop point-to-point, mesh point-to-point, or point-to-multipoint LSP template

lsp-name

the name that identifies the LSP

Output

The following output is an example of RSVP LSP egress statistical information.

Output example

```
a.show>router>mpls>lsp-egress-stats lsp_1
-----
LSP Name      : toNodeC_1
-----
Collect Stats : Disabled          Accting Plcy. : None
Adm State     : Up                PSB Match    : True
FC BE
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC L2
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC AF
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC L1
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC H2
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC EF
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC H1
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC NC
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
-----
LSP Egress Statistics : 1
=====
```

lsp-ingress-stats

Syntax

lsp-ingress-stats [**type** *lsp-type*] [**active**] [**template-match** *SessionNameString* [**sender** *ip-address*]]
lsp-ingress-stats **lsp** *lsp-name* **sender** *ip-address*

Context

show>router>mpls

Description

This command displays RSVP LSP ingress statistical information.

Parameters

lsp-type

specifies the type of LSP to display. The only available options are **p2p** and **p2mp**.

active

displays information from all LSPs with statistics collection enabled

template-match

displays information for a one-hop point-to-point, mesh point-to-point, or point-to-multipoint LSP template

SessionNameString

the name of the session

ip-address

the system IP address of the sender

lsp-name

the name that identifies the LSP

Output

The following output is an example of RSVP LSP ingress statistical information.

Output example

```
b.show>router>mpls>lsp-ingress-stats 1.2.2.2 lsp lsp_1
=====
MPLS LSPs Ingress Statistics
=====
-----
LSP Name       : toNodeA_1
Sender         : 10.10.10.29
-----
Collect Stats  : Disabled           Accting Plcy. : None
Adm State      : Up                 PSB Match    : True
FC BE
InProf Pkts   : 0                   OutProf Pkts : 0
InProf Octets : 0                   OutProf Octets: 0
FC L2
InProf Pkts   : 0                   OutProf Pkts : 0
InProf Octets : 0                   OutProf Octets: 0
FC AF
InProf Pkts   : 0                   OutProf Pkts : 0
InProf Octets : 0                   OutProf Octets: 0
FC L1
InProf Pkts   : 0                   OutProf Pkts : 0
InProf Octets : 0                   OutProf Octets: 0
FC H2
InProf Pkts   : 0                   OutProf Pkts : 0
InProf Octets : 0                   OutProf Octets: 0
FC EF
InProf Pkts   : 0                   OutProf Pkts : 0
InProf Octets : 0                   OutProf Octets: 0
FC H1
InProf Pkts   : 0                   OutProf Pkts : 0
InProf Octets : 0                   OutProf Octets: 0
FC NC
InProf Pkts   : 0                   OutProf Pkts : 0
InProf Octets : 0                   OutProf Octets: 0
```

```
-----
LDP Ingress Statistics : 1
=====
```

lsp-template

Syntax

lsp-template [*template-name*] **bindings**
lsp-template [*template-name*] [**detail**]

Context

show>router>mpls

Description

This command displays MPLS LSP template information.

Parameters

- template-name*
the unique name for the LSP template
- bindings**
displays any bindings associated with the LSP template
- detail**
displays detailed information for the LSP template

Output

The following output is an example of MPLS LSP template information, and [Table 17: Router MPLS LSP template field descriptions](#) describes the fields.

Output example

```
A:ALU-12# show router mpls lsp-template detail
=====
MPLS LSP Templates (Detail)
=====
-----
LSP Template : MeshTemplateWithLoosePath
-----
Type           : MeshP2P           Admin State    : Up
From           : 10.20.1.3
Default Path   : LooseHopPathNameW* Adaptive        : Enabled
Bandwidth      : 0 Mbps             Hop Limit      : 18
CSPF           : Enabled            Use TE metric   : Disabled
Propagate Admin Grp: Enabled
Include Groups :                    Exclude Groups  :
None                                     G0
                                       G1
                                       G2
                                       G3
FastReroute    : Enabled
FR Method      : Facility           FR Hop Limit   : 13
```

```

FR Prop Admin Group: Enabled
FR Node Protect      : Disabled
Record Route         : Record
Retry Limit          : 100
LSP Count            : 3
LdpOverRsvp          : Disabled
IGP Shortcut         : Enabled
IGP LFA              : Disabled
Least Fill           : Enabled
SetupPriority        : 7
Egress Stats         : Disabled
Collect Stats        : Disabled
Class Type           : 0
Main CT Retry Limit  : 0
ADSPEC               : Disabled
Record Label         : Record
Retry Timer          : 30 sec
Ref Count            : 0
VprnAutoBind         : Enabled
BGP Shortcut         : Disabled
IGP Rel Metric       : Disabled
Metric               : 25
Hold Priority        : 0
Accounting Policy    : None
Backup Class Type    : 0
BGP Transport Tunn   : Enabled
=====
A:ALU-12#

```

Table 17: Router MPLS LSP template field descriptions

Label	Description
Type	The type of LSP
Admin State	Down – the path is administratively disabled
	Up – the path is administratively enabled
From	The IP address of the ingress router for the LSP
Default Path	The value used to order the hops in a path
Adaptive	Indicates whether the adaptive option is enabled or disabled
Bandwidth	n/a
Hop Limit	The maximum number of hops that an LSP can traverse, including the ingress and egress routers
CSPF	Indicates whether the CSPF option is enabled or disabled (always enabled for LSP templates)
Use TE metric	Indicates whether the TE metric option is enabled or disabled
Propagate Admin Grp	Indicates whether the propagate admin group option is enabled or disabled
Include Groups	The admin groups that are to be included by an LSP when signaling a path
Exclude Groups	The admin groups that are to be excluded by an LSP when signaling a path
FastReroute	Indicates whether the FRR option is enabled or disabled
FR Method	The type of FRR that is used by the path (always Facility for LSP templates)

Label	Description
FR Hop Limit	The total number of hops a protection LSP can take before merging back onto the main LSP path
FR Prop Admin Group	Indicates whether the FRR propagate admin group option is enabled or disabled
FR Node Protect	Indicates whether FRR has node protection enabled or disabled
Record Route	Indicates whether the route is being recorded
Record Label	Indicates whether the label is being recorded
Retry Limit	The maximum number of retries allowed
Retry Timer	The time between retry attempts
LSP Count	The number of LSPs belonging to the LSP template
Ref Count	n/a
LdpOverRsvp	n/a
VprnAutoBind	Indicates whether the VPRN autobind option is enabled or disabled
IGP Shortcut	Indicates whether the IGP shortcut option is enabled or disabled
BGP Shortcut	n/a
IGP LFA	Indicates whether the IGP LFA option is enabled or disabled
IGP Rel Metric	Indicates whether the IGP relative metric option is enabled or disabled
Least Fill	Indicates whether the least fill option is enabled or disabled
Metric	The TE metric value
SetupPriority	n/a
Hold Priority	n/a
Egress Stats	n/a
Collect Stats	n/a
Accounting Policy	n/a
Class Type	n/a
Backup Class Type	n/a

Label	Description
Main CT Retry Limit	n/a
BGP Transport Tunn	Indicates whether the BGP transport tunnel option is enabled or disabled
ADSPEC	enabled – the LSP will include advertising data (ADSPEC) objects in RSVP-TE messages
	disabled – the LSP will not include advertising data (ADSPEC) objects in RSVP-TE messages

path

Syntax

path [path-name] [lsp-binding]

Context

show>router>mpls

Description

This command displays MPLS paths.

Parameters

path-name
the unique name label for the LSP path

lsp-binding
displays binding information

Output

The following output is an example of MPLS path information, and [Table 18: Router MPLS path field descriptions](#) describes the fields.

Output example

```
A:ALU-12# show router mpls path
=====
MPLS Path:
=====
```

Path Name	Adm	Hop Index	IP Address	Strict/Loose
nyc_to_sjc_via_dfw	Up	20	192.20.1.4	Strict
		30	192.20.1.6	Strict
		40	192.20.1.8	Strict
		50	192.20.1.10	Strict
nyc_to_sjc_via_den	Up	10	192.20.1.5	Strict
		20	192.20.1.7	Loose
		30	192.20.1.9	Loose

```

                                40      192.20.1.11      Loose
                                50      192.20.1.13      Strict
secondary_path2                Down no hops      n/a      n/a
-----
Paths : 3
=====
A:ALU-12#

A:ALU-12# show router mpls path lsp-binding
=====
MPLS Path:
=====
Path Name                      Opr  LSP Name                      Binding
-----
nyc_to_sjc_via_dfw             Up   NYC_SJC_customer1            Primary
nyc_to_sjc_via_den             Up   NYC_SJC_customer1            Standby
secondary_path2                Down NYC_SJC_customer1            Seconda*
-----
Paths : 3
=====
A:ALU-12#

```

Table 18: Router MPLS path field descriptions

Label	Description
Path Name	The unique name label for the LSP path
Adm	Down – the path is administratively disabled
	Up – the path is administratively enabled
Hop Index	The value used to order the hops in a path
IP Address	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 other routers
Opr	The operational status of the path (up or down)
LSP Name	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	Total number of paths configured

sr-te-lsp

Syntax

sr-te-lsp [*lsp-name*] [**status** {**up** | **down**}] [**detail**] **path** [*path-name*]

sr-te-lsp [*lsp-name*] [**detail**]

sr-te-lsp [*sp-name*] [**status** {**up** | **down**}] [**to ip-address**] [**detail**]

Context

show>router>mpls

Description

This command displays segment routing-traffic engineering (SR-TE) LSP information.

Parameters

lsp-name

displays SR-TE LSPs with the specified LSP name

status

displays SR-TE LSPs with the specified status (up or down)

detail

displays detailed information for the LSP template

path-name

displays SR-TE LSPs with the specified path name

ip-address

displays SR-TE LSPs connected to the IP address

Output

The following outputs are examples of MPLS SR-TE path information.

- [Output example](#)
- [Output example: PCEP SR-TE LSP configurations](#)

Output example

```
*A:Dut-C# show router mpls sr-te-lsp detail

=====
MPLS SR-TE LSPs (Originating) (Detail)
=====
Legend :
+ - Inherited
=====
-----
Type : Originating
-----
LSP Name   : to_A_SR_TE
LSP Type   : SrTeLsp
LSP Index  : 65536
LSP Tunnel ID : 1
TTM Tunnel Id : 655362
```

```

From      : 3.3.3.3
To        : 1.1.1.1
Adm State : Up
LSP Up Time : 0d 00:00:05
Transitions : 1
Retry Limit : 0
Hop Limit   : 255
PathCompMethod : none
FallbkPathComp : not-applicable
Metric      : N/A
Local Sr Protec*: preferred
Load Bal Wt : N/A
Include Grps :
None
Egress Stats : Disabled
BFD Template : None
BFD Enable   : False
WaitForUpTimer : 4
ReturnPathLabel : None

Oper State : Up
LSP Down Time : 0d 00:00:00
Path Changes  : 1
Retry Timer   : 2 sec
Negotiated MTU : 1496

Label Stack Reduction: Disabled
ClassForwarding      : Disabled
Exclude Grps         :
None

BFD Ping Intvl : N/A
BFD Failure-action : None

Revert Timer : Disabled
Entropy Label : Enabled+
Negotiated EL : Disabled
VprnAutoBind : Enabled
IGP Shortcut  : Enabled
IGP LFA       : Disabled
BGPTransTun   : Enabled
Oper Metric   : 16777215
PCE Report    : Disabled+
PCE Control   : Disabled
Max SR Labels : 6
Path Profile  : None
Admin Tags    : None

Next Revert In : N/A
Oper Entropy Label : Enabled
Override Tunnel ELC : Disabled

BGP Shortcut : Enabled
IGP Rel Metric : Disabled

Additional FRR Labels: 1

Primary(a) : path_A
Up Time : 0d 00:00:09

Bandwidth : 0 Mbps
Secondary : path_A2

Down Time : 0d 00:00:50

Bandwidth : 0 Mbps
-----
Type : Originating
-----
LSP Name : template-1.1.1.1-671747
LSP Type : OneHopP2PSrTe
LSP Index : 81921
LSP Tunnel ID : 16386
From : 30.30.2.2
To : 1.1.1.1
TTM Tunnel Id : 671747
Adm State : Up
LSP Up Time : 0d 00:00:08
Oper State : Up
Transitions : 1
LSP Down Time : 0d 00:00:00
Retry Limit : 0
Path Changes : 1
PathCompMethod : none
Retry Timer : 30 sec
FallbkPathComp : not-applicable
Metric : N/A
Local Sr Protec*: preferred
Label Stack Reduction: Disabled
Load Bal Wt : N/A
ClassForwarding : Disabled
Include Grps :
Exclude Grps :
None
Egress Stats : Disabled
BFD Template : srte
BFD Ping Intvl : N/A
BFD Enable : True
BFD Failure-action : None
WaitForUpTimer : 4
ReturnPathLabel : 35002

```

```

Revert Timer      : Disabled          Next Revert In      : N/A
Entropy Label    : Enabled+          Oper Entropy Label  : Enabled
Negotiated EL    : Disabled          Override Tunnel ELC : Disabled
VprnAutoBind     : Enabled
IGP Shortcut     : Enabled           BGP Shortcut       : Enabled
IGP LFA          : Disabled          IGP Rel Metric     : Disabled
BGPTransTun     : Enabled
Oper Metric      : 16777215
OriginTemplName  : template

Admin Tags       : None
PCE Report       : Disabled+
PCE Control      : Disabled
Max SR Labels    : 6                 Additional FRR Labels: 1
Path Profile     : None

Primary(a)       : path_A            Up Time             : 0d 00:00:09
Bandwidth        : 0 Mbps
=====
* indicates that the corresponding row element may have been truncated.

```

```
*A:Dut-C# show router mpls sr-te-lsp path detail
```

```

=====
MPLS SR-TE LSP
Path (Detail)
=====
Legend :
  S      - Strict                      L      - Loose
  A-SID  - Adjacency SID              N-SID  - Node SID
  +      - Inherited
=====
-----
LSP SR-TE to_A_SR_TE
Path path_A
-----
LSP Name      : to_A_SR_TE
Path LSP ID   : 3072
From          : 3.3.3.3
To            : 1.1.1.1
Admin State   : Up                    Oper State     : Up
Path Name     : path_A
Path Type     : Primary
Path Admin    : Up                    Path Oper      : Up
Path Up Time  : 0d 00:01:11           Path Down Time : 0d 00:00:00
Retry Limit   : 0                     Retry Timer    : 2 sec
Retry Attempt : 0                     Next Retry In  : 0 sec

PathCompMethod : none                 OperPathCompMethod: none
MetricType     : igp                  Oper MetricType  : igp
LocalSrProt    : preferred            Oper LocalSrProt : N/A
LabelStackRed  : Disabled             Oper LabelStackRed: N/A

Bandwidth      : No Reservation        Oper Bandwidth   : 0 Mbps
Hop Limit      : 255                   Oper HopLimit    : 255
Setup Priority  : 7                     Oper SetupPriority: 7
Hold Priority   : 0                     Oper HoldPriority : 0
Inter-area     : N/A

PCE Updt ID    : 0                     PCE Updt State  : None
PCE Upd Fail Code: noError

```

PCE Report	: Disabled+	Oper PCE Report	: Disabled
PCE Control	: Disabled	Oper PCE Control	: Disabled
Include Groups	:	Oper IncludeGroups:	
None		None	
Exclude Groups	:	Oper ExcludeGroups:	
None		None	
Last Resignal	: n/a		
IGP/TE Metric	: 16777215	Oper Metric	: 16777215
Oper MTU	: 1496	Path Trans	: 1
Degraded	: False		
Failure Code	: noError		
Failure Node	: n/a		
Explicit Hops	:		
No Hops Specified			
Actual Hops	:		
1.1.1.1(1.1.1.1)(N-SID)		Record Label	: 19410
BFD Configuration and State			
Template	: srte	Ping Interval	: N/A
Enable	: True	State	: up
ReturnPathLabel	: 35001		
WaitForUpTimer	: 4 sec	OperWaitForUpTimer:	4 sec
WaitForUpTmLeft	: 0		
StartFail Rsn	: N/A		

LSP SR-TE to_A_SR_TE			
Path path_A2			

LSP Name	: to_A_SR_TE		
Path LSP ID	: 3074		
From	: 3.3.3.3		
To	: 1.1.1.1		
Admin State	: Up	Oper State	: Up
Path Name	: path_A2		
Path Type	: Secondary		
Path Admin	: Up	Path Oper	: Down
Path Up Time	: 0d 00:00:00	Path Down Time	: 0d 00:01:53
Retry Limit	: 0	Retry Timer	: 2 sec
Retry Attempt	: 0	Next Retry In	: 0 sec
PathCompMethod	: none	OperPathCompMethod:	N/A
MetricType	: igp	Oper MetricType	: N/A
LocalSrProt	: preferred	Oper LocalSrProt	: N/A
LabelStackRed	: Disabled	Oper LabelStackRed:	N/A
Bandwidth	: No Reservation	Oper Bandwidth	: N/A
Hop Limit	: 255	Oper HopLimit	: N/A
Setup Priority	: 7	Oper SetupPriority:	N/A
Hold Priority	: 0	Oper HoldPriority	: N/A
Inter-area	: N/A		
PCE Updt ID	: 0	PCE Updt State	: None
PCE Upd Fail Code:	noError		
PCE Report	: Disabled+	Oper PCE Report	: Disabled
PCE Control	: Disabled	Oper PCE Control	: Disabled
Include Groups	:	Oper IncludeGroups:	
None		N/A	
Exclude Groups	:	Oper ExcludeGroups:	

```

None
Last Resignal      : n/a

IGP/TE Metric      : N/A
Oper MTU           : N/A
Degraded           : False
Failure Code       : noError
Failure Node       : n/a
Explicit Hops      :
  No Hops Specified
Actual Hops        :
  No Hops Specified
Srlg               : Disabled

Srlg Disjoint      : False

BFD Configuration and State
Template           : srte
Enable            : True
ReturnPathLabel    : 35002
WaitForUpTimer     : 4 sec
WaitForUpTmLeft    : 0
StartFail Rsn      : N/A

Ping Interval      : N/A
State              : notApplicable

OperWaitForUpTimer: 4 sec

-----
LSP SR-TE template-1.1.1.1-671747
Path path_A
-----
LSP Name          : template-1.1.1.1-671747
Path LSP ID       : 14336
From              : 30.30.2.2
To                : 1.1.1.1
Admin State       : Up
Oper State        : Up
Path Name         : path_A
Path Type         : Primary
Path Admin        : Up
Path Up Time      : 0d 00:01:12
Path Oper         : Up
Path Down Time    : 0d 00:00:00
Retry Limit       : 0
Retry Timer       : 30 sec
Retry Attempt     : 0
Next Retry In     : 0 sec

PathCompMethod    : none
MetricType        : igp
LocalSrProt       : preferred
LabelStackRed     : Disabled
OperPathCompMethod: none
Oper MetricType   : igp
Oper LocalSrProt  : N/A
Oper LabelStackRed: N/A

Bandwidth         : No Reservation
Hop Limit         : 2
Setup Priority     : 7
Hold Priority      : 0
Inter-area        : N/A
Oper Bandwidth    : 0 Mbps
Oper HopLimit     : 2
Oper SetupPriority: 7
Oper HoldPriority  : 0

PCE Updt ID       : 0
PCE Upd Fail Code: noError
PCE Updt State    : None

PCE Report        : Disabled+
PCE Control       : Disabled
Oper PCE Report   : Disabled
Oper PCE Control  : Disabled

Include Groups    :
None
Exclude Groups   :
None
Last Resignal     : n/a
Oper IncludeGroups:
None
Oper ExcludeGroups:
None

IGP/TE Metric      : 16777215
Oper Metric        : 16777215
Oper MTU           : 1496
Path Trans        : 1
Degraded           : False

```



```

Failure Code      : noError
Failure Node      : n/a
Explicit Hops     :
  No Hops Specified
Actual Hops       :
  30.30.2.1(1.1.1.1) (A-SID)
Record Label      : 524282

BFD Configuration and State
Template          : None
Enable            : False
ReturnPathLabel   : None
WaitForUpTimer    : 4 sec
WaitForUpTmLeft   : 0
StartFail Rsn     : N/A
Ping Interval     : N/A
State             : up
OperWaitForUpTimer: 4 sec

```

```
*A:7705:Dut-A# show router mpls sr-te-lsp status up
```

```
=====
MPLS SR-TE LSPs (Originating)
=====
```

LSP Name	To	Tun Id	Protect Path	Adm	Opr
sr-te-lsp-to-Dut-C	10.20.1.3	1	N/A	Up	Up
sr-te-lsp-to-Dut-B	10.20.1.22	3	N/A	Up	Up
sr-te-lsp-to-Dut-D	10.20.1.4	4	N/A	Up	Up

```

-----
LSPs : 3
=====

```

```
*A:7705:Dut-A# show router mpls sr-te-lsp to 10.20.1.4
```

```
=====
MPLS SR-TE LSPs (Originating)
=====
```

LSP Name	To	Tun Id	Protect Path	Adm	Opr
sr-te-lsp-to-Dut-D	10.20.1.4	4	N/A	Up	Up

```

-----
LSPs : 1
=====

```

```
*A:7705:Dut-A# show router mpls sr-te-lsp to 10.20.1.4 detail
```

```
=====
MPLS SR-TE LSPs (Originating) (Detail)
=====
```

```
-----
Type : Originating
-----
```

LSP Name	: sr-te-lsp-to-Dut-D	LSP Tunnel ID	: 4
LSP Type	: SrTeLsp	TTM Tunnel Id	: 655365
LSP Index	: 65539	To	: 10.20.1.4
From	: 10.20.1.1	Oper State	: Up
Adm State	: Up	LSP Down Time	: 0d 00:00:00
LSP Up Time	: 0d 00:20:22	Path Changes	: 1
Transitions	: 1	Retry Timer	: 20 sec
Retry Limit	: 0	Negotiated MTU	: 1542
Hop Limit	: 255		
CSPF	: Disabled	Exclude Grps	:
Metric	: N/A	None	
Include Grps	:		
None			
VprnAutoBind	: Enabled		

```

IGP Shortcut      : Enabled
IGP LFA           : Disabled
BGPTransTun      : Enabled
Oper Metric       : 16777215
PCE Report        : Inherited
PCE Compute       : Disabled
Max SR Labels     : 4
Path Profile      :
None
Primary(a)        : sr-te-to-Dut-D
Bandwidth         : 0 Mbps
=====
IGP Rel Metric    : Disabled
PCE Control       : Disabled
Additional FRR Labels: 1
Up Time           : 0d 00:20:22
=====

```

```

*A:7705:Dut-A# show router mpls sr-te-lsp status up detail path "sr-te-lsp-to-Dut-D"
=====

```

```

MPLS SR-TE LSP sr-te-lsp-to-Dut-D Path (Detail)
=====

```

```

Legend :

```

```

S - Strict

```

```

L - Loose
=====
-----

```

```

SR-TE LSP sr-te-lsp-to-Dut-D Path sr-te-to-Dut-D
-----

```

```

LSP Name          : sr-te-lsp-to-Dut-D
Path LSP ID       : 57856
From              : 10.20.1.1
Admin State       : Up
Path Name         : sr-te-to-Dut-D
Path Admin        : Up
Path Up Time      : 0d 00:22:23
Retry Limit       : 0
Retry Attempt     : 2
CSPF              : Disabled
Bandwidth         : No Reservation
Hop Limit         : 255
Setup Priority     : 7
Hold Priority      : 0
Inter-area        : N/A

To                : 10.20.1.4
Oper State        : Up
Path Type         : Primary
Path Oper         : Up
Path Down Time    : 0d 00:00:00
Retry Timer       : 20 sec
Next Retry In     : 0 sec
Oper CSPF         : Disabled
Oper Bandwidth    : 0 Mbps
Oper HopLimit     : 255
Oper Setup Priority : 7
Oper Hold Priority : 0

PCE Updt ID       : 0
PCE Upd Fail Code : noError
PCE Report        : Inherited
PCE Control       : Disabled
PCE Compute       : Disabled
Include Groups    :
None
Exclude Groups    :
None

PCE Updt State    : None
Oper PCE Report   : Disabled
Oper PCE Control  : Disabled
Oper PCE Compute  : Disabled
Oper Include Groups :
None
Oper Exclude Groups :
None

IGP/TE Metric     : 16777215
Oper MTU          : 1542
Failure Code       : noError
Failure Node       : n/a
Explicit Hops      :
435287
-> 435087
Actual Hops        :
n/a
-> n/a

Oper Metric        : 16777215
Path Trans         : 1

Record Label       : 435287
Record Label       : 435287

```

```
=====
*A:7705:Dut-A# show router mpls sr-te-lsp "sr-te-lsp-to-Dut-D"
=====
```

```
MPLS SR-TE LSPs (Originating)
=====
```

LSP Name	To	Tun Id	Protect Path	Adm	Opr
sr-te-lsp-to-Dut-D	10.20.1.4	4	N/A	Up	Up

```
-----
LSPs : 1
=====
```

```
=====
*A:7705:Dut-A# show router mpls sr-te-lsp "sr-te-lsp-to-Dut-D" path detail
=====
```

```
MPLS SR-TE LSP sr-te-lsp-to-Dut-D Path (Detail)
=====
```

```
Legend :
```

```
S - Strict
```

```
L - Loose
=====
```

```
-----
SR-TE LSP sr-te-lsp-to-Dut-D Path sr-te-to-Dut-D
-----
```

LSP Name	: sr-te-lsp-to-Dut-D	
Path LSP ID	: 57856	
From	: 10.20.1.1	To : 10.20.1.4
Admin State	: Up	Oper State : Up
Path Name	: sr-te-to-Dut-D	Path Type : Primary
Path Admin	: Up	Path Oper : Up
Path Up Time	: 0d 00:23:15	Path Down Time : 0d 00:00:00
Retry Limit	: 0	Retry Timer : 20 sec
Retry Attempt	: 2	Next Retry In : 0 sec
CSPF	: Disabled	Oper CSPF : Disabled
Bandwidth	: No Reservation	Oper Bandwidth : 0 Mbps
Hop Limit	: 255	Oper HopLimit : 255
Setup Priority	: 7	Oper Setup Priority : 7
Hold Priority	: 0	Oper Hold Priority : 0
Inter-area	: N/A	
PCE Updt ID	: 0	PCE Updt State : None
PCE Upd Fail Code	: noError	
PCE Report	: Inherited	Oper PCE Report : Disabled
PCE Control	: Disabled	Oper PCE Control : Disabled
PCE Compute	: Disabled	Oper PCE Compute : Disabled
Include Groups	:	Oper Include Groups :
None		None
Exclude Groups	:	Oper Exclude Groups :
None		None
IGP/TE Metric	: 16777215	Oper Metric : 16777215
Oper MTU	: 1542	Path Trans : 1
Failure Code	: noError	
Failure Node	: n/a	
Explicit Hops	:	
435287		
-> 435087		
Actual Hops	:	
n/a		Record Label : 435287
-> n/a		Record Label : 435287

```
=====
```

```
*A:7705:Dut-A#
```

Output example: PCEP SR-TE LSP configurations

The following CLI displays are output examples of a PCEP SR-TE LSP in three configurations:

- PCE-computation enabled, PCE-report disabled (via inheritance), and PCE-control disabled
- PCE-computation enabled, PCE-report enabled, and PCE-control disabled
- PCE-computation enabled, PCE-report enabled, and PCE-control enabled

The configuration can be determined by checking the PathCompMethod, PCE Report, and PCE Control fields.

Output example

```
*A:7705:Dut-C# show router mpls sr-te-lsp detail
=====
MPLS SR-TE LSPs (Originating) (Detail)
=====
Type : Originating
-----
LSP Name       : test_lsp_1
LSP Type       : SrTeLsp
LSP Index      : 65536
From           : 10.20.1.3
Adm State      : Up
LSP Up Time    : 0d 00:00:44
Transitions    : 1
Retry Limit    : 0
Hop Limit      : 255
CSPF           : Enabled
Metric         : N/A
Include Grps   :
None           :

LSP Tunnel ID  : 1
TTM Tunnel Id  : 655362
To             : 10.20.1.4
Oper State     : Up
LSP Down Time  : 0d 00:00:00
Path Changes   : 1
Retry Timer    : 30 sec
Negotiated MTU : 1492

Use TE metric  : Disabled
Exclude Grps   :
None           :

VprnAutoBind   : Enabled
IGP Shortcut   : Enabled
IGP LFA        : Disabled
BGPTransTun    : Enabled
Oper Metric    : 100
PCE Report     : Inherited
PathCompMethod : pce
Max SR Labels  : 6
Path Profile    :
None           :

IGP Rel Metric : Disabled

PCE Control    : Disabled
Additional FRR Labels: 1

Primary(a)     : fully_loose
Bandwidth      : 0 Mbps
Up Time        : 0d 00:00:44
=====
*A:7705:Dut-C#
```

```
*A:7705:Dut-C# show router mpls sr-te-lsp detail
=====
MPLS SR-TE LSPs (Originating) (Detail)
=====
Type : Originating
-----
LSP Name       : test_lsp_1
LSP Type       : SrTeLsp
LSP Index      : 65536
From           : 10.20.1.3
LSP Tunnel ID  : 1
TTM Tunnel Id  : 655362
To             : 10.20.1.4
```

```

Adm State       : Up
LSP Up Time     : 0d 00:00:02
Transitions     : 3
Retry Limit     : 0
Hop Limit       : 255
CSPF            : Enabled
Metric          : N/A
Include Grps    :
None
VprnAutoBind    : Enabled
IGP Shortcut    : Enabled
IGP LFA         : Disabled
BGPTransTun     : Enabled
Oper Metric     : 100
PCE Report      : Enabled
PathCompMethod  : pce
Max SR Labels   : 6
Path Profile    :
None

Oper State      : Up
LSP Down Time   : 0d 00:00:00
Path Changes    : 3
Retry Timer     : 30 sec
Negotiated MTU  : 1492

Use TE metric   : Disabled
Exclude Grps    :
None

IGP Rel Metric  : Disabled

PCE Control     : Disabled
Additional FRR Labels: 1

Primary(a)      : fully_loose
Bandwidth       : 0 Mbps
=====
*A:7705:Dut-C#

```

```

*A:7705:Dut-C# show router mpls sr-te-lsp detail
=====
MPLS SR-TE LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name       : test_lsp_1
LSP Type       : SrTeLsp
LSP Index      : 65536
From           : 10.20.1.3
Adm State      : Up
LSP Up Time    : 0d 00:00:42
Transitions    : 3
Retry Limit    : 0
Hop Limit      : 255
CSPF           : Enabled
Metric         : N/A
Include Grps   :
None

LSP Tunnel ID  : 1
TTM Tunnel Id  : 655362
To             : 10.20.1.4
Oper State     : Up
LSP Down Time  : 0d 00:00:00
Path Changes   : 3
Retry Timer    : 30 sec
Negotiated MTU : 1492

Use TE metric  : Disabled
Exclude Grps   :
None

VprnAutoBind   : Enabled
IGP Shortcut    : Enabled
IGP LFA        : Disabled
BGPTransTun    : Enabled
Oper Metric    : 100
PCE Report     : Enabled
PathCompMethod : pce
Max SR Labels  : 6
Path Profile   :
None

IGP Rel Metric : Disabled

PCE Control    : Enabled
Additional FRR Labels: 1

Primary(a)     : fully_loose
Bandwidth      : 0 Mbps
Up Time        : 0d 00:00:42
=====
*A:7705:Dut-C#

```

static-lsp

Syntax

```
static-lsp [/sp-name]
static-lsp [/sp-type]
static-lsp count
```

Context

```
show>router>mpls
```

Description

This command displays MPLS static LSP information.

Parameters

- lsp-name*
name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique.
- lsp-type*
type that identifies the LSP. The LSP type is one of the keywords **transit** or **terminate**, where **terminate** displays the number of static LSPs that terminate at the router, and **transit** displays the number of static LSPs that transit the router.
- count**
the number of static LSPs that originate and terminate at the router

Output

The following output is an example of MPLS static LSP information, and [Table 19: Router MPLS static LSP field descriptions](#) describes the fields.

Output example

```
ALU-12# show router mpls static-lsp
=====
MPLS Static LSPs (Originating)
=====
LSP Name      To          Next Hop      Out Label Up/Down Time  Adm  Opr
ID            ID          ID            ID
-----
to131         10.9.9.9    10.1.2.2      131      30d 02:42:53  Up   Down
1            n/a
to121         10.8.8.8    10.1.3.2      121      30d 02:42:53  Up   Down
2            n/a
static-lsp_   10.9.9.9    10.1.2.2      35        0d 01:39:34  Up   Down
cc
3            n/a
-----
LSPs : 3
=====
*A:ALU-12>show>router>mpls#
```

Output example - static-lsp transit

```
A:ALU-12# show router mpls static-lsp transit
=====
MPLS Static LSPs (Transit)
=====
In Label    In I/F      Out Label   Out I/F     Next Hop          Adm  Opr
-----
1020        1/1/1      1021        1/1/5       10.10.10.6        Up   Up
-----
LSPs : 1
=====
```

Output example - static-lsp terminate

```
*A:ALU-12>show>router>mpls# static-lsp terminate
=====
MPLS Static LSPs (Terminate)
=====
In Label    In Port     Out Label   Out Port    Next Hop          Adm  Opr
-----
131         1/3/1      n/a         n/a         n/a               Up   Down
121         1/2/1      n/a         n/a         n/a               Up   Down
35          1/3/1      n/a         n/a         n/a               Up   Down
-----
LSPs : 3
=====
```

Output example - static-lsp count

```
*A:ALU-12>show>router>mpls# static-lsp count
=====
MPLS Static-LSP Count
=====
Originate      Transit      Terminate
-----
0              0            0
=====
*A:ALU-12>show>router>mpls# static-lsp
```

Table 19: Router MPLS static LSP field descriptions

Label	Description
Lsp Name	The name of the LSP used in the path
To	The system IP address of the egress router for the LSP
Next Hop	The system IP address of the next hop in the LSP path
Out Label	The egress label
Adm	Down – indicates that the path is administratively disabled
	Up – indicates that the path is administratively enabled
Opr	Down – indicates that the path is operationally down

Label	Description
	Up – indicates that the path is operationally up
LSPs	The total number of static LSPs
In Label	The ingress label
In Port	The ingress port
Out Port	The egress port
Up/Down Time	The duration that the LSP is either operationally up or down
Static-LSP Count	The number of originating, transit, and terminating static LSPs

status

Syntax

status

Context

show>router>mpls

Description

This command displays MPLS operation information.

Output

The following output is an example of MPLS status information, and [Table 20: Router MPLS status field descriptions](#) describes the fields.

Output example

```
A:NOK-Dut-B# show router mpls status
=====
MPLS Status
=====
Admin Status           : Up           Oper Status           : Up
Oper Down Reason       : n/a
FRR Object              : Enabled   Resignal Timer        : Disabled
Hold Timer              : 1 seconds Next Resignal         : N/A
Admin Group Frr         : Disabled
Dynamic Bypass          : Enabled   User Srlg Database    : Disabled
BypassResignalTimer     : Disabled BypassNextResignal    : N/A
LeastFill Min Thd       : 5 percent LeastFill Reopti Thd  : 10 percent
Local TTL Prop          : Enabled   Transit TTL Prop      : Enabled
Exp Backoff Retry       : Disabled CSPF On Loose Hop     : Disabled
Lsp Init RetryTimeout   : 30 seconds
Logger Event Bundling   : Disabled
RetryIgpOverload        : Disabled
Sec FastRetryTimer      : Disabled Static LSP FR Timer    : 30 seconds
P2PActPathFastRetry     : Disabled
In Maintenance Mode     : No
```



```

Pce-report          : None
Next Available Lsp Index : 4
Entropy Label RSVP-TE : Enabled      Entropy Label SR-TE      : Enabled
=====
MPLS LSP Count
=====
-----
                Originate          Transit          Terminate
-----
Static LSPs      0                  0              0
Dynamic LSPs     0                  0              0
Detour LSPs      0                  0              0
Mesh-P2P LSPs   0                  N/A            N/A
One Hop-P2P LSPs 0                  N/A            N/A
SR-TE LSPs      0                  N/A            N/A
=====
A:NOK-Dut-B#

```

Table 20: Router MPLS status field descriptions

Label	Description
Admin Status	Down – indicates that MPLS is administratively disabled
	Up – indicates that MPLS is administratively enabled
Oper Status	Down – indicates that MPLS is operationally down
	Up – indicates that MPLS is operationally up
FRR 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	The amount of time that the ingress node holds before programming its data plane and declaring the LSP up to the service module
Oper Down Reason	The reason that MPLS is operationally down
Next Resignal	The amount of time until the next resignal for the LSP
Dynamic Bypass	Indicates whether dynamic bypass is enabled or disabled
Next Available Lsp Index	The next free LSP index ID
Entropy Label RSVP-TE	Enabled – specifies that entropy label is enabled for the RSVP-TE LSP

Label	Description
	Disabled – specifies that entropy label is disabled for the RSVP-TE LSP
Entropy Label SR-TE	Enabled – specifies that entropy label is enabled for the SR-TE LSP
	Disabled – specifies that entropy label is disabled for the SR-TE LSP
MPLS LSP Counts	The number of originate, transit, and terminate LSPs that are Static, Dynamic, Detour, Mesh-P2P, One Hop-P2P, or SR-TE

3.24.2.4 Show commands (MPLS-labels)

label

Syntax

label start-label [end-label | [in-use |label-owner]]

Context

show>router>mpls-labels

Description

This command displays 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, evpn, ildp, mirror, rsvp, static, sr, svcmgr, tldp, vprn

Output

The following output is an example of MPLS label information, and [Table 21: Router MPLS-labels label field descriptions](#) describes the fields.

Output example

```
ALU-12# show router mpls-labels label 131070
=====
MPLS Label 131070
=====
Label                Label Type          Label Owner
-----
131070               dynamic             RESERVED-POOL
-----
In-use labels in entire range : 71103
=====
ALU-12#
```

```
*A: Sar18 Dut-B>show>router>mpls-labels># label 32 131071 ildp
=====
MPLS Labels from 32 to 131071 (Owner: ILDP)
=====
Label                Label Type          Label Owner
-----
131070               dynamic             ILDP
131071               dynamic             ILDP
-----
In-use labels (Owner: ILDP) in specified range : 2
In-use labels (Owner: All) in specified range : 2
In-use labels in entire range : 2
=====
```

Table 21: Router MPLS-labels label field descriptions

Label	Description
Label	The value of the label
Label Type	Specifies whether the label value is statically or dynamically assigned
Label Owner	The label owner
In-use labels (Owner: ILDP) in specified range	The total number of labels in the specified range being used by the specified owner
In-use labels (Owner: All) in specified range	The total number of labels in the specified range being used by all owners
In-use labels in entire range	The total number of labels being used

label-range

Syntax

label-range

Context

show>router>mpls-labels

Description

This command displays the MPLS label range.

Output

The following output is an example of MPLS label range information, and [Table 22: Router MPLS-labels label range field descriptions](#) describes the fields.

Output example

```
ALU-12# show router mpls-labels label-range
=====
Label Ranges
=====
Label Type      Start Label End Label  Aging      Available  Total
-----
Static          32          18431      -          18400     18400
Dynamic         18432       131071     0          41537     112640
  Seg-Route     20000       90000      -           0         70001
-----
Reserved Label Blocks
-----
Reserved Label      Start      End      Total
Block Name          Label      Label
-----
Block100            120001     120100     100
a123456789012345678901234567890123 130070     131071     1002
45678901234567891111
-----
No. of Reserved Label Blocks: 2
-----
=====
ALU-12#
```

Table 22: Router MPLS-labels label range field descriptions

Label	Description
Label Type	Displays information about static-lsp , static-svc , and dynamic label types
Start Label	The label value assigned at the ingress router
End Label	The label value assigned for the egress router

Label	Description
Aging	The number of labels released from a service that are transitioning back to the label pool. Labels are aged for 15 seconds.
Available	The number of label values available in the label range
Total	The total number of label values in the label range
Reserved label Block Name	The name of the reserved label block
Start Label	The starting value of the label range in the reserved label block
End Label	The end value of the label range in the reserved label block
No. of Reserved Label Blocks	The number of reserved label blocks configured on the system

summary

Syntax

summary

Context

show>router>mpls-labels

Description

This command displays a summary of MPLS label usage.

Output

The following output is an example of MPLS label summary information, and [Table 23: Router MPLS-labels summary field descriptions](#) describes the fields.

Output example

```
*A:Sar18 Dut-B>show>router>mpls-labels># summary
=====
Mpls-Labels Summary
=====
Static Label Range           : 18400
Bgp Labels Hold Timer        : 0
Segment Routing Start Label   : 0
Segment Routing End Label     : 0
Reserved Label Block Name    :
    Block100
    a12345678901234567890123456789012345678901234567891111
=====
```

Table 23: Router MPLS-labels summary field descriptions

Label	Description
Static Label Range	The number of label values available in the static label range
Bgp Labels Hold Timer	The number of BGP labels in use under control of the hold timer
Segment Routing Start Label	The start label value for segment routing labels
Segment Routing End Label	The end label value for segment routing labels
Reserved Label Block Name	The name of the block for a range of labels that are reserved for local assignment

3.24.2.5 Show commands (seamless BFD)

session

Syntax

session lsp-index *lsp-index* [**path-lspid** *path-lspid*] [**detail**]

session lsp-name *lsp-name* [**path-lspid** *path-lspid*] [**detail**]

session lsp-path [**detail**]

session lsp-path prefix *ip-prefix/prefix-length* [**src** *ip-address*] [**detail**]

Context

show>router>bfd>seamless-bfd

Description

This command displays details about seamless BFD (S-BFD) sessions.

The **session lsp-index** command displays S-BFD session details based on an LSP index and optional LSP ID.

The **session lsp-name** command displays S-BFD session details based on an LSP name and optional LSP ID.

The **session lsp-path [detail]** command displays a summary of all the S-BFD sessions that are bound to LSPs.

The **session lsp-path prefix** *ip-prefix/prefix-length* [**src** *ip-address*] [**detail**] command displays information about an S-BFD session to a specified far-end prefix.

Parameters

- lsp-index*
the LSP index value, from 0 to 4294967295
- path-lspid*
the LSP path identifier, from 0 to 4294967295
- lsp-name*
the LSP path name
- ip-prefix/prefix-length*
the IP prefix and prefix length associated with the LSP path
- ip-address*
the source IP address

Output

The following show output displays an LSP that has BFD configured to use a controlled return path. [Table 24: MPLS seamless BFD session field descriptions](#) describes the fields.

Output example

```
*A:Dut-C># show router bfd seamless-bfd session lsp-path detail

=====
BFD Session
=====
Prefix      : 10.20.1.6/32
Local Address : 10.20.1.3
LSP Name    : C_F_70001
LSP Index   : 70001                Path LSP ID      : 14336
Fec Type    : srTe
Return Path : 454600
Oper State  : Up                   Protocols       : mplsLsp
Up Time     : 3d 02:43:10          Up Transitions  : 1
Down Time   : None                Down Transitions : 0
                                           Version Mismatch : 0

Forwarding Information
Local Discr  : 8                   Local State     : Up
Local Diag  : 0 (None)
Local Mode  : Demand
Local Min Tx : 1000
Last Sent (ms) : 0                Local Mult      : 3
Type        : cpm-np              Local Min Rx    : 1000
Remote Discr : 0                   Remote State    : Up
Remote Diag  : 0 (None)            Remote Mode     : Demand
Remote Min Tx : 1000               Remote Mult     : 3
Remote C-flag : 1                  Remote Min Rx   : 1000
Last Recv (ms) : 0
=====
=====
```

Table 24: MPLS seamless BFD session field descriptions

Label	Description
Prefix	Specifies the prefix

Label	Description
Local Address	Specifies the local address
LSP Name	Specifies the LSP name
LSP Index	Specifies the LSP index
Path LSP ID	Specifies the path LSP ID
Fec Type	Specifies the FEC type
Return Path	Specifies the return path label for the S-BFD packet
Oper State	Specifies the operational state (Up/Down)
Protocols	Specifies the protocol used
Up Time	Specifies the up time
Up Transitions	Specifies the number of up transitions
Down Time	Specifies the down time
Down Transitions	Specifies the number of down transitions
Version Mismatch	Specifies the number of version mismatches
Local Discr	Specifies the local Discr information
Remote State	Specifies the remote state (Up/Down)
Local Diag	Specifies the local diag information
Local Mode	Specifies the local mode
Local Min Tx	Specifies the local Min Tx value
Local Mult	Specifies the local Mult value
Last Sent (ms)	Specifies the last sent forwarding information
Local Min Rx	Specifies the local Min Rx value
Type	Specifies the session endpoint type
Remote Discr	Specifies the remote Discr value
Remote State	Specifies the remote state (Up/Down)
Remote Diag	Specifies the remote diagnostic value
Remote Mode	Specifies the remote mode
Remote Min Tx	Specifies the remote minimum transmit value
Remote Mult	Specifies the remote multiplier
Remote C-flag	Specifies the remote C-flag value

Label	Description
Last Recv (ms)	Specifies the last received value
Remote Min Rx	Specifies the remote minimum receive value

3.24.2.6 Show commands (RSVP)



Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

interface

Syntax

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

Context

show>router>rsvp

Description

This command shows RSVP-TE interface information.

Parameters

- ip-int-name*
identifies the network IP interface. The interface name cannot be in the form of an IP address. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes
- ip-address*
the system or network interface IP address
- statistics**
the IP address and the number of packets sent and received on an per-interface basis
- detail**
displays detailed information

Output

The following outputs are examples of RSVP-TE interface information:

- RSVP-TE Interface ([Output example](#), [Table 25: Router RSVP-TE interface field descriptions](#))
- RSVP-TE Interface Detail ([Output example](#), [Table 26: Router RSVP-TE interface detail field descriptions](#))
- RSVP-TE Interface Statistics ([Output example](#), [Table 27: Router RSVP-TE interface statistics field descriptions](#))

Output example

```

A:ALU-12# show router rsvp interface
=====
RSVP Interfaces
=====
Interface                Total    Active    Total BW    Resv BW    Adm Opr
                        Sessions Sessions (Mbps)    (Mbps)
-----
system                   -         -         -         -         Up  Up
ip-10.10.1.1             1         1        100         0         Up  Up
ip-10.10.2.1             1         1        100         0         Up  Up
ip-10.10.3.1             0         0        100         0         Up  Up
-----
Interfaces : 4
=====
A:ALU-12#

```

Table 25: Router RSVP-TE interface field descriptions

Label	Description
Interface	The name of the IP interface
Total Sessions	The total number of RSVP-TE sessions on this interface. This count includes sessions that are active as well as sessions that have been signaled but a response has not yet been received.
Active Sessions	The total number of active RSVP-TE sessions on this interface
Total BW (Mbps)	The amount of bandwidth in megabits per second (Mbps) available to be reserved for the RSVP-TE protocol on the interface
Resv BW (Mbps)	The amount of bandwidth in megabits per second (Mbps) reserved on this interface. A value of zero (0) indicates that no bandwidth is reserved.
Adm	Down – the RSVP-TE interface is administratively disabled
	Up – the RSVP-TE interface is administratively enabled
Opr	Down – the RSVP-TE interface is operationally down
	Up – the RSVP-TE interface is operationally up
Interfaces	The number of interfaces listed in the display

Output example

```

A: ALU-12# show router rsvp interface "ip-10.10.1.1" detail
=====
RSVP Interfaces (Detailed): ip-10.10.1.1
-----
Interface : ip-10.10.1.1
-----
Interface      : ip-10.10.1.1

```

```

Port ID       : 1/1/1
Admin State   : Up
Active Sessions : 0
Total Sessions : 0
Subscription  : 10 %
Total BW      : 100 Mbps
Hello Interval : 3000 ms
Key Type Auth : Disabled
Keychain Auth : Disabled
Auth Rx Seq Num : n/a
Auth Tx Seq Num : n/a
Refresh Reduc. : Disabled
Bfd Enabled    : No
ImplicitNullLabel : Disabled*

Oper State    : Up
Active Resvs  : 0
Port Speed    : 1000 Mbps
Aggregate     : Dsabl
Hello Timeouts : 0

Auth Key Id   : n/a
Auth Win Size : n/a
Reliable Deli. : Disabled
Graceful Shut. : Disabled
GR helper     : Disabled

IGP Update
Up Thresholds(%) : 0 15 30 45 60 75 80 85 90 95 96 97 98 99 100 *
Down Thresholds(%) : 100 99 98 97 96 95 90 85 80 75 60 45 30 15 0 *
IGP Update Pending: No
Next Update      : N/A
No Neighbors.
* Indicates Inherited Values
-----
A: ALU-12#

```

Table 26: Router RSVP-TE interface detail field descriptions

Label	Description
Interface	The name of the IP interface
Port ID	The physical port bound to the interface
Admin State	Down – the RSVP-TE interface is administratively disabled
	Up – the RSVP-TE interface is administratively enabled
Oper State	Down – the RSVP-TE interface is operationally down
	Up – the RSVP-TE interface is operationally up
Active Sessions	The total number of active RSVP-TE sessions on this interface
Active Resvs	The total number of active RSVP-TE sessions that have reserved bandwidth
Total Sessions	The total number of RSVP-TE sessions on this interface. This count includes sessions that are active as well as sessions that have been signaled but a response has not yet been received.
Subscription	The percentage of the link bandwidth that RSVP-TE can use for reservation. When the value is zero (0), no new sessions are permitted on this interface.
Port Speed	The speed on the interface

Label	Description
Total BW	The amount of bandwidth in megabits per second (Mbps) available to be reserved for the RSVP-TE protocol on this interface
Aggregate	Indicates whether aggregate messages are sent. Aggregate messages are used to pack multiple RSVP messages into a single packet to reduce the network overhead. When the value is true, RSVP negotiates with each neighbor and gets consensus before sending aggregate messages.
Hello Interval	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. A value of zero (0) indicates that the sending of hello messages is disabled. A value of n/a indicates that the interface is an unnumbered interface.
Hello Timeouts	The total number of Hello messages that timed out on this RSVP-TE interface. A value of n/a indicates that the interface is an unnumbered interface.
Auth Rx Seq Num	The received MD5 sequence number
Auth Key Id	The MD5 key identifier
Auth Tx Seq Num	The transmitted MD5 sequence number
Auth Win Size	The MD5 window size
Refresh Reduc.	Indicates whether refresh reduction capabilities are enabled or disabled
Reliable Deli.	Indicates whether reliable delivery is enabled or disabled
Bfd Enabled	Indicates whether BFD is enabled or disabled on the RSVP-TE interface. A value of n/a indicates that BFD is not applicable because the interface is an unnumbered interface.
Graceful Shut.	Indicates whether graceful shutdown is enabled or disabled
ImplicitNullLabel	Indicates whether the implicit null label is enabled or disabled
GR helper	Indicates whether Graceful-Restart Helper is enabled or disabled
IGP Update	
Up Thresholds (%)	Indicates up threshold levels for the interface
Down Thresholds (%)	Indicates down threshold levels for the interface
IGP Update Pending	Indicates whether an IGP update will occur

Label	Description
Next Update	Indicates when the next IGP update will be, if there is one pending

Output example

```

A:ALU-12# show router rsvp interface statistics
=====
RSVP Interface (statistics)
=====
Interface system
-----
Interface           : Up
Total Packets       (Sent) : 0                (Recd.): 0
Bad Packets         (Sent) : 0                (Recd.): 0
Paths               (Sent) : 0                (Recd.): 0
Path Errors         (Sent) : 0                (Recd.): 0
Path Tears          (Sent) : 0                (Recd.): 0
Resvs               (Sent) : 0                (Recd.): 0
Resv Confirms       (Sent) : 0                (Recd.): 0
Resv Errors         (Sent) : 0                (Recd.): 0
Resv Tears          (Sent) : 0                (Recd.): 0
Refresh Summaries   (Sent) : 0                (Recd.): 0
Refresh Acks        (Sent) : 0                (Recd.): 0
Bundle Packets      (Sent) : 0                (Recd.): 0
Hellos              (Sent) : 0                (Recd.): 0
Auth Errors         (Sent) : 0                (Recd.): 0
-----

```

Table 27: Router RSVP-TE interface statistics field descriptions

Label	Description
Interface	The name of the IP interface displayed in the header
Interface (status)	The status of the interface (up or down)
Sent	The total number of error-free RSVP-TE packets that have been transmitted on the RSVP-TE interface
Recd	The total number of error-free RSVP-TE packets received on the RSVP-TE interface
Total Packets	The total number of RSVP-TE packets, including errors, received on the RSVP-TE interface
Bad Packets	The total number of RSVP-TE packets with errors transmitted on the RSVP-TE interface
Paths	The total number of RSVP-TE PATH messages received on the RSVP-TE interface
Path Errors	The total number of RSVP-TE PATH ERROR messages transmitted on the RSVP-TE interface

Label	Description
Path Tears	The total number of RSVP-TE PATH TEAR messages received on the RSVP-TE interface
Resvs	The total number of RSVP-TE RESV messages received on the RSVP-TE interface
Resv Confirms	The total number of RSVP-TE RESV CONFIRM messages received on the RSVP-TE interface
Resv Errors	The total number of RSVP-TE RESV ERROR messages received on the RSVP-TE interface
Resv Tears	The total number of RSVP-TE RESV TEAR messages received on the RSVP-TE interface
Refresh Summaries	The total number of RSVP-TE RESV summary refresh messages received on the RSVP-TE interface
Refresh Acks	The total number of RSVP-TE RESV acknowledgment messages received when refresh reduction is enabled on the RSVP-TE interface
Bundle Packets	The total number of RSVP-TE RESV bundle packets received on the RSVP-TE interface
Hellos	The total number of RSVP-TE RESV HELLO REQ messages received on the RSVP-TE interface
Auth Errors	The number of authentication errors

neighbor

Syntax

neighbor [*ip-address*][**detail**]

Context

show>router>rsvp

Description

This command displays RSVP-TE neighbors.

Parameters

ip-address

the IP address of the originating router

detail

displays detailed information

Output

The following output is an example of RSVP-TE neighbor information, and [Table 28: Router RSVP-TE neighbor field descriptions](#) describes the fields.

Output example

```
*A:ALU-12>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
=====
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
=====
```

Table 28: Router RSVP-TE neighbor field descriptions

Label	Description
Neighbor	The IP address of the RSVP-TE neighbor
Interface	The interface ID of the RSVP-TE neighbor
Hello	The status of the Hello message
Last Oper Change	The time of the last operational change to the connection
Flags	Any flags 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

Description

This command shows RSVP-TE session information.

Parameters

session-type
specifies the session type

Values originate, transit, terminate, detour, detour-transit, detour-terminate, bypass-tunnel, manual-bypass

from ip-address

specifies the IP address of the originating router

to ip-address

specifies the IP address of the egress router

name

specifies the name of the LSP used in the path

status up

specifies to display a session that is operationally up

status down

specifies to display a session that is operationally down

detail

displays detailed information

Output

The following output is an example of RSVP-TE session information, and [Table 29: Router RSVP-TE session field descriptions](#) describes the fields.

Output example

```
A:ALU-12# show router rsvp session
=====
RSVP Sessions
=====
```

From	To	Tunnel ID	LSP ID	Name	State
10.20.1.3	10.20.1.1	1	37	C_A_1::C_A_1	Up
10.20.1.3	10.20.1.1	2	38	C_A_2::C_A_2	Up
10.20.1.3	10.20.1.1	3	39	C_A_3::C_A_3	Up
10.20.1.3	10.20.1.1	4	40	C_A_4::C_A_4	Up
10.20.1.1	10.20.1.3	2	40	A_C_2::A_C_2	Up
10.20.1.1	10.20.1.3	3	41	A_C_3::A_C_3	Up
10.20.1.1	10.20.1.3	4	42	A_C_4::A_C_4	Up
10.20.1.1	10.20.1.3	5	43	A_C_5::A_C_5	Up
10.20.1.1	10.20.1.3	6	44	A_C_6::A_C_6	Up
10.20.1.1	10.20.1.3	7	45	A_C_7::A_C_7	Up
10.20.1.1	10.20.1.3	8	46	A_C_8::A_C_8	Up
10.20.1.3	10.20.1.1	5	41	C_A_5::C_A_5	Up
10.20.1.3	10.20.1.1	6	42	C_A_6::C_A_6	Up
10.20.1.3	10.20.1.1	7	43	C_A_7::C_A_7	Up
10.20.1.3	10.20.1.1	8	44	C_A_8::C_A_8	Up
...					

```
-----
Sessions : 65
=====
A:ALU-12#
```

```
A:ALU-12# show router rsvp session lsp-name A_C_2::A_C_2 status up
=====
RSVP Sessions
```


From	To	Tunnel ID	LSP ID	Name	State
10.20.1.1	10.20.1.3	2	40	A_C_2::A_C_2	Up
Sessions : 1					
A:ALU-12#					

Table 29: Router RSVP-TE session field descriptions

Label	Description
From	The IP address of the originating router
To	The IP address of the egress router
Tunnel ID	The ID of the ingress node of the tunnel supporting this RSVP-TE session
LSP ID	The ID assigned by the agent to this RSVP-TE session
Name	The administrative name assigned to the RSVP-TE session by the agent
State	Down – the operational state of this RSVP-TE session is down
	Up – the operational state of this RSVP-TE session is up

statistics

Syntax
statistics

Context
show>router>rsvp

Description
This command displays global statistics in the RSVP-TE instance.

Output
The following output is an example of RSVP-TE statistics information, and [Table 30: Router RSVP-TE statistics field descriptions](#) describes the fields.

Output example

```
A:ALU-12# show router rsvp statistics
=====
RSVP Global Statistics
=====
```

```

PATH Timeouts      : 0
GR Helper PATH Tim*: 0
RESV Timeouts      : 0
GR Helper RESV Tim*: 0
=====

```

Table 30: Router RSVP-TE statistics field descriptions

Label	Description
PATH Timeouts	The total number of PATH timeouts
RESV Timeouts	The total number of RESV timeouts
GR Helper PATH Timeouts	The total number of graceful restart helper PATH timeouts
GR Helper RESV Timeouts	The total number of graceful restart helper RESV timeouts

status

Syntax

status

Context

show>router>rsvp

Description

This command displays RSVP-TE operational status.

Output

The following output is an example of RSVP-TE status information, and [Table 31: Router RSVP-TE status field descriptions](#) describes the fields.

Output example

```

A:ALU-12# show router rsvp status
=====
RSVP Status
=====
Admin Status      : Down      Oper Status      : Down
Keep Multiplier   : 3         Refresh Time     : 30 sec
Message Pacing    : Disabled   Pacing Period    : 100 msec
Max Packet Burst  : 650 msgs   Refresh Bypass   : Disabled
Rapid Retransmit  : 5 hmsec    Rapid Retry Limit : 3
Graceful Shutdown : Disabled   SoftPreemptionTimer: 300 sec
GR Max Recovery   : 300 sec    GR Max Restart   : 120 sec
Implicit Null Label: Disabled   Node-id in RR0   : Exclude
P2P Merge Point Ab*: Disabled  P2MP Merge Point A*: Disabled
DiffServTE AdmModel: Basic     Entropy Label    : Disabled
Percent Link Bw CT0: 100       Percent Link Bw CT4: 0
Percent Link Bw CT1: 0         Percent Link Bw CT5: 0
Percent Link Bw CT2: 0         Percent Link Bw CT6: 0
Percent Link Bw CT3: 0         Percent Link Bw CT7: 0
TE0 -> Class Type : 0         Priority          : 0
TE1 -> Class Type : 0         Priority          : 1

```

```

TE2 -> Class Type : 0          Priority : 2
TE3 -> Class Type : 0          Priority : 3
TE4 -> Class Type : 0          Priority : 4
TE5 -> Class Type : 0          Priority : 5
TE6 -> Class Type : 0          Priority : 6
TE7 -> Class Type : 0          Priority : 7
IgpThresholdUpdate : Disabled
Up Thresholds(%) : 0 15 30 45 60 75 80 85 90 95 96 97 98 99 100
Down Thresholds(%) : 100 99 98 97 96 95 90 85 80 75 60 45 30 15 0
Update Timer : N/A
Update on CAC Fail : Disabled
=====

```

Table 31: Router RSVP-TE status field descriptions

Label	Description
Admin Status	Down – RSVP-TE is administratively disabled
	Up – RSVP-TE is administratively enabled
Oper Status	Down – RSVP-TE is operationally down
	Up – RSVP-TE is operationally up
Keep Multiplier	The keep-multiplier number used by RSVP-TE to declare that a reservation is down or the neighbor is down
Refresh Time	The refresh-time interval , in seconds, between the successive PATH and RESV refresh messages
Message Pacing	Enabled – RSVP-TE messages, specified in the max-burst command, are sent in a configured interval, specified in the period command
	Disabled – message pacing is disabled. RSVP-TE message transmission is not regulated.
Pacing Period	The time interval, in milliseconds, during which the router can send the number of RSVP-TE messages specified in the max-burst command
Max Packet Burst	The maximum number of RSVP-TE messages that are sent under normal operating conditions in the period specified
Refresh Bypass	Enabled – the refresh-reduction-over-bypass command is enabled
	Disabled – the refresh-reduction-over-bypass command is disabled
Rapid Retransmit	The time interval for the rapid retransmission time, which is used in the retransmission mechanism that handles unacknowledged message_id objects (the units "hmsec")

Label	Description
	represent hundreds of msec; for example, 5 hmsec represents 500 msec)
Rapid Retry Limit	The value of the rapid retry limit, which is used in the retransmission mechanism that handles unacknowledged message_id objects
Graceful Shutdown	Specifies whether graceful shutdown of the RSVP node is enabled

3.24.2.7 Clear commands

interface

Syntax

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

Context

clear>router>mpls

Description

This command resets or clears statistics for MPLS interfaces.

Parameters

ip-int-name

specifies an existing IP interface. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

statistics

clears only statistics

lsp

Syntax

lsp [*lsp-name*]

Context

clear>router>mpls

Description

This command resets and restarts an LSP.

Parameters

lsp-name

specifies the name of the LSP to clear

lsp-egress-stats

Syntax

lsp-egress-stats [*lsp-name*]

Context

clear>router>mpls

Description

This command clears RSVP LSP egress statistics.



Note: When RSVP LSP statistics are cleared, the current aggregate statistics count is recorded as a baseline and is used to provide a relative count each time the statistics are viewed with the **show** command. Because this baseline number is not reconciled between the active and inactive CSMs, after a CSM activity switch the statistics on the newly active CSM will show the aggregate count as though no clear command has been executed.

Parameters

lsp-name

specifies the name of the LSP to clear

lsp-ingress-stats

Syntax

lsp-ingress-stats [*ip-address* **lsp** *lsp-name*]

Context

clear>router>mpls

Description

This command clears RSVP LSP ingress statistics.



Note: When RSVP LSP statistics are cleared, the current aggregate statistics count is recorded as a baseline and is used to provide a relative count each time the statistics are viewed with the **show** command. Because this baseline number is not reconciled between the active and inactive CSMs, after a CSM activity switch the statistics on the newly active CSM will show the aggregate count as though no clear command has been executed.

Parameters

ip-address

the system IP address of the sender

lsp-name

the name that identifies the LSP

interface

Syntax

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

Context

clear>router>rsvp

Description

This command resets or clears statistics for an RSVP-TE interface.

Parameters

ip-int-name

identifies the IP interface to clear. The interface name cannot be in the form of an IP address. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes

statistics

clears only statistics

statistics

Syntax

statistics

Context

clear>router>rsvp

Description

This command clears global statistics for the RSVP-TE instance; for example, clears **path** and **resv timeout** counters.

3.24.2.8 Monitor commands

interface

Syntax

interface *interface* [*interface...*[(up to 5 max)] [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]

Context

monitor>router>mpls
monitor>router>rsvp

Description

This command displays statistics for MPLS or RSVP interfaces at the configured interval until the configured count is reached.

The first screen displays the current statistics related to the MPLS or RSVP interface. The subsequent statistical information listed for each interval is displayed as a delta to the previous display. When the keyword **rate** is specified, the rate-per-second for each statistic is displayed instead of the delta.

Monitor commands are similar to **show** commands but only statistical information displays. Monitor commands display the selected statistics according to the configured number of times at the interval specified.

Default

n/a

Parameters

interface
specifies the name of the IP interface or the IP address

Values *ip-int-name* | *ip-address*

seconds
configures the interval for each display in seconds

Values 11 to 60 (MPLS)
 3 to 60 (RSVP)

Default 11 (MPLS)
 10 (RSVP)

repeat
configures how many times the command is repeated

Values 1 to 999

Default 10

absolute

displays raw statistics, without processing. No calculations are performed on the delta or rate statistics.

rate

displays the rate per second for each statistic instead of the delta

lsp-egress-stats

Syntax

lsp-egress-stats **lsp** *lsp-name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]

Context

monitor>router>mpls

Description

This command monitors MPLS LSP egress statistics at the configured interval until the configured count is reached.

Default

no lsp-egress-stats

Parameters

lsp-name

the LSP name

seconds

configures the interval for each display in seconds

Values 3 to 60

Default 10

repeat

configures how many times the command is repeated

Values 1 to 999

Default 10

absolute

displays raw statistics, without processing. No calculations are performed on the delta or rate statistics.

rate

displays the rate per second for each statistic instead of the delta

Output

The following output is an example of LSP egress statistical information.

Output example

```
=====
Monitor ingress statistics for MPLS LSP "toNodeA_1"
=====
-----
At time t = 0 sec (Base Statistics)
-----
LSP Name      : toNodeA_1
Sender        : 10.10.10.29
-----
Collect Stats : Disabled          Accting Plcy. : None
Adm State     : Up                PSB Match    : True
FC BE
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC L2
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC AF
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC L1
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC H2
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC EF
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC H1
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC NC
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
-----
At time t = 3 sec (Mode: Absolute)
-----
LSP Name      : toNodeA_1
Sender        : 10.10.10.29
-----
Collect Stats : Disabled          Accting Plcy. : None
Adm State     : Up                PSB Match    : True
FC BE
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC L2
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC AF
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC L1
InProf Pkts   : 0                 OutProf Pkts : 0
InProf Octets : 0                 OutProf Octets: 0
FC H2
```

```

InProf Pkts   : 0
InProf Octets : 0
FC EF
InProf Pkts   : 0
InProf Octets : 0
FC H1
InProf Pkts   : 0
InProf Octets : 0
FC NC
InProf Pkts   : 0
InProf Octets : 0
OutProf Pkts  : 0
OutProf Octets: 0
-----
At time t = 6 sec (Mode: Absolute)
-----
LSP Name      : toNodeA_1
Sender        : 10.10.10.29
-----
Collect Stats : Disabled
Adm State     : Up
FC BE
InProf Pkts   : 0
InProf Octets : 0
FC L2
InProf Pkts   : 0
InProf Octets : 0
FC AF
InProf Pkts   : 0
InProf Octets : 0
FC L1
InProf Pkts   : 0
InProf Octets : 0
FC H2
InProf Pkts   : 0
InProf Octets : 0
FC EF
InProf Pkts   : 0
InProf Octets : 0
FC H1
InProf Pkts   : 0
InProf Octets : 0
FC NC
InProf Pkts   : 0
InProf Octets : 0
Accting Plcy. : None
PSB Match     : True
OutProf Pkts  : 0
OutProf Octets: 0
OutProf Pkts  : 0
OutProf Octets: 0
OutProf Pkts  : 0
OutProf Octets: 0
OutProf Pkts  : 0
OutProf Octets: 0
OutProf Pkts  : 0
OutProf Octets: 0
OutProf Pkts  : 0
OutProf Octets: 0
OutProf Pkts  : 0
OutProf Octets: 0
-----
At time t = 9 sec (Mode: Absolute)
-----

```

lsp-ingress-stats

Syntax

lsp-ingress-stats **lsp** *lsp-name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] *ip-address*

Context

monitor>router>mpls

Description

This command displays MPLS LSP ingress statistics at the configured interval until the configured count is reached.

Default

no lsp-egress-stats

Parameters

lsp-name

the LSP name

seconds

configures the interval for each display in seconds

Values 3 to 60

Default 10

repeat

configures how many times the command is repeated

Values 1 to 999

Default 10

absolute

displays raw statistics, without processing. No calculations are performed on the delta or rate statistics.

rate

displays the rate per second for each statistic instead of the delta

ip-address

the IP address of the remote host

Output

The following output is an example of LSP ingress statistical information.

Output example

```
B:Sys# monitor router mpls lsp-ingress-stats lsp sample 1.1.1.1 repeat 3
interval 10 absolute
=====
Monitor ingress statistics for MPLS LSP "sample"
-----
At time t = 0 sec (Base Statistics)
-----
LSP Name : sample
Sender : 1.1.1.1
-----
Collect Stats : Enabled Accting Plcy. : None
Adm State : Up PSB Match : True
FC BE
InProf Pkts : 539 OutProf Pkts : 0
```

```
InProf Octets : 548702 OutProf Octets: 0
FC L2
InProf Pkts : 0 OutProf Pkts : 539
InProf Octets : 0 OutProf Octets: 548702
FC AF
InProf Pkts : 0 OutProf Pkts : 0
InProf Octets : 0 OutProf Octets: 0
FC L1
InProf Pkts : 1078 OutProf Pkts : 0
InProf Octets : 1097404 OutProf Octets: 0
FC H2
InProf Pkts : 0 OutProf Pkts : 539
InProf Octets : 0 OutProf Octets: 548702
FC EF
InProf Pkts : 539 OutProf Pkts : 0
InProf Octets : 548702 OutProf Octets: 0
FC H1
InProf Pkts : 539 OutProf Pkts : 0
InProf Octets : 548702 OutProf Octets: 0
FC NC
InProf Pkts : 0 OutProf Pkts : 539
InProf Octets : 0 OutProf Octets: 548702
```

At time t = 10 sec (Mode: Absolute)

LSP Name : sample
Sender : 1.1.1.1

Collect Stats : Enabled Accting Plcy. : None

Adm State : Up PSB Match : True

FC BE

InProf Pkts : 568 OutProf Pkts : 0

InProf Octets : 578224 OutProf Octets: 0

FC L2

InProf Pkts : 0 OutProf Pkts : 568

InProf Octets : 0 OutProf Octets: 578224

FC AF

InProf Pkts : 0 OutProf Pkts : 0

InProf Octets : 0 OutProf Octets: 0

FC L1

InProf Pkts : 1136 OutProf Pkts : 0

InProf Octets : 1156448 OutProf Octets: 0

FC H2

InProf Pkts : 0 OutProf Pkts : 568

InProf Octets : 0 OutProf Octets: 578224

FC EF

InProf Pkts : 568 OutProf Pkts : 0

InProf Octets : 578224 OutProf Octets: 0

FC H1

InProf Pkts : 568 OutProf Pkts : 0

InProf Octets : 578224 OutProf Octets: 0

FC NC

InProf Pkts : 0 OutProf Pkts : 568

InProf Octets : 0 OutProf Octets: 578224

At time t = 20 sec (Mode: Absolute)

LSP Name : sample

Sender : 1.1.1.1

Collect Stats : Enabled Accting Plcy. : None

Adm State : Up PSB Match : True

FC BE

InProf Pkts : 597 OutProf Pkts : 0

```
InProf Octets : 607746 OutProf Octets: 0
FC L2
InProf Pkts : 0 OutProf Pkts : 597
InProf Octets : 0 OutProf Octets: 607746
FC AF
InProf Pkts : 0 OutProf Pkts : 0
InProf Octets : 0 OutProf Octets: 0
FC L1
InProf Pkts : 1194 OutProf Pkts : 0
InProf Octets : 1215492 OutProf Octets: 0
FC H2
InProf Pkts : 0 OutProf Pkts : 597
InProf Octets : 0 OutProf Octets: 607746
FC EF
InProf Pkts : 597 OutProf Pkts : 0
InProf Octets : 607746 OutProf Octets: 0
FC H1
InProf Pkts : 597 OutProf Pkts : 0
InProf Octets : 607746 OutProf Octets: 0
FC NC
InProf Pkts : 0 OutProf Pkts : 597
InProf Octets : 0 OutProf Octets: 607746
-----
At time t = 30 sec (Mode: Absolute)
-----
LSP Name : sample
Sender : 1.1.1.1
-----
Collect Stats : Enabled Accting Plcy. : None
Adm State : Up PSB Match : True
FC BE
InProf Pkts : 627 OutProf Pkts : 0
InProf Octets : 638286 OutProf Octets: 0
FC L2
InProf Pkts : 0 OutProf Pkts : 627
InProf Octets : 0 OutProf Octets: 638286
FC AF
InProf Pkts : 0 OutProf Pkts : 0
InProf Octets : 0 OutProf Octets: 0
FC L1
InProf Pkts : 1254 OutProf Pkts : 0
InProf Octets : 1276572 OutProf Octets: 0
FC H2
InProf Pkts : 0 OutProf Pkts : 627
InProf Octets : 0 OutProf Octets: 638286
FC EF
InProf Pkts : 627 OutProf Pkts : 0
InProf Octets : 638286 OutProf Octets: 0
FC H1
InProf Pkts : 627 OutProf Pkts : 0
InProf Octets : 638286 OutProf Octets: 0
FC NC
InProf Pkts : 0 OutProf Pkts : 627
InProf Octets : 0 OutProf Octets: 638286
=====
B:Sys#
```

3.24.2.9 Debug commands

mpls

Syntax

[no] mpls [**lsp** *lsp-name*] [**sender** *source-address*] [**endpoint** *endpoint-address*] [**tunnel-id** *tunnel-id*] [**lsp-id** *lsp-id*] [**interface** *ip-int-name*]

Context

debug>router

Description

This command enables and configures debugging for MPLS.

Parameters

lsp-name

the name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique.

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

ip-int-name

identifies the interface. The interface name cannot be in the form of an IP address. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

event

Syntax

[no] event

Context

```
debug>router>mpls
```

```
debug>router>rsvp
```

Description

This command enables debugging for specific events.

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

```
all
```

Syntax

```
all [detail]
```

```
no all
```

Context

```
debug>router>mpls>event
```

```
debug>router>rsvp>event
```

Description

This command debugs all events.

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

Parameters

```
detail
```

displays detailed information about all events

```
frr
```

Syntax

```
frr [detail]
```

```
no frr
```

Context

```
debug>router>mpls>event
```

Description

This command debugs fast reroute events.

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

Parameters**detail**

displays detailed information about reroute events

iom**Syntax**

iom [**detail**]

no iom

Context

debug>router>mpls>event

Description

This command debugs MPLS IOM events.

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

Parameters**detail**

displays detailed information about MPLS IOM events

lsp-setup**Syntax**

lsp-setup [**detail**]

no lsp-setup

Context

debug>router>mpls>event

Description

This command debugs LSP setup events.

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

Parameters**detail**

displays detailed information about LSP setup events

mbb

Syntax

mbb [detail]

no mbb

Context

debug>router>mpls>event

Description

This command debugs the state of the most recent invocation of the make-before-break (MBB) functionality.

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

Parameters

detail

displays detailed information about MBB events

misc

Syntax

misc [detail]

no misc

Context

debug>router>mpls>event

debug>router>rsvp>event

Description

This command debugs miscellaneous events.

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

Parameters

detail

displays detailed information about miscellaneous events

XC

Syntax

xc [**detail**]

no xc

Context

debug>router>mpls>event

Description

This command debugs cross-connect events.

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

Parameters

detail

displays detailed information about cross-connect events

rsvp

Syntax

[**no**] **rsvp** [**lsp** *lsp-name*] [**sender** *sender-address*] [**endpoint** *endpoint-address*] [**tunnel-id** *tunnel-id*] [**lsp-id** *lsp-id*] [**interface** *ip-int-name*]

no rsvp

Context

debug>router

Description

This command enables and configures debugging for RSVP.

Parameters

lsp-name

name that identifies the LSP

sender-address

specifies the system IP address of the sender

endpoint-address

specifies the far-end system IP address

tunnel-id

specifies the RSVP-TE tunnel ID

Values 0 to 4294967295

lsp-id

specifies the LSP ID

Values 1 to 65535

ip-int-name

identifies the interface

auth

Syntax

auth

no auth

Context

debug>router>rsvp>event

Description

This command debugs authentication events.

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

Parameters

detail

displays detailed information about authentication events

nbr

Syntax

nbr [detail]

no nbr

Context

debug>router>rsvp>event

Description

This command debugs neighbor events.

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

Parameters

detail

displays detailed information about neighbor events

path

Syntax

path [**detail**]

no path

Context

debug>router>rsvp>event

Description

This command debugs path-related events.

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

Parameters

detail

displays detailed information about path-related events

resv

Syntax

resv [**detail**]

no resv

Context

debug>router>rsvp>event

Description

This command debugs RSVP-TE reservation events.

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

Parameters

detail

displays detailed information about RSVP-TE reservation events

```
rr
```

Syntax

```
rr  
no rr
```

Context

```
debug>router>rsvp>event
```

Description

This command debugs refresh reduction events.

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

Parameters

```
detail
```

displays detailed information about refresh reduction events

```
packet
```

Syntax

```
[no] packet
```

Context

```
debug>router>rsvp
```

Description

This command enters the context to debug packets.

```
ack
```

Syntax

```
ack [detail]  
no ack
```

Context

```
debug>router>rsvp>packet
```

Description

This command debugs ack packets.

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

Parameters

detail

displays detailed information about RSVP-TE ack packets

all

Syntax

all [detail]

no all

Context

debug>router>rsvp>packet

Description

This command debugs all packets.

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

Parameters

detail

displays detailed information about all RSVP-TE packets

bundle

Syntax

bundle [detail]

no bundle

Context

debug>router>rsvp>packet

Description

This command debugs bundle packets.

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

Parameters

detail

displays detailed information about RSVP-TE bundle packets

hello

Syntax

hello [detail]

no hello

Context

debug>router>rsvp>packet

Description

This command debugs hello packets.

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

Parameters

detail

displays detailed information about hello packets

path

Syntax

path [detail]

no path

Context

debug>router>rsvp>packet

Description

This command enables debugging for RSVP-TE path packets.

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

Parameters

detail

displays detailed information about path-related events

patherr

Syntax

patherr [detail]

no patherr

Context

debug>router>rsvp>packet

Description

This command debugs path error packets.

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

Parameters

detail

displays detailed information about path error packets

pathtear

Syntax

pathtear [detail]

no pathtear

Context

debug>router>rsvp>packet

Description

This command debugs path tear packets.

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

Parameters

detail

displays detailed information about path tear packets

resv

Syntax

resv [detail]

no resv

Context

debug>router>rsvp>packet

Description

This command enables debugging for RSVP-TE RESV packets.

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

Parameters

detail

displays detailed information about RSVP-TE RESV packets

```
resvterr
```

Syntax

resvterr [detail]

no resvterr

Context

debug>router>rsvp>packet

Description

This command debugs ResvErr packets.

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

Parameters

detail

displays detailed information about ResvErr packets

```
resvtear
```

Syntax

resvtear [detail]

no resvtear

Context

debug>router>rsvp>packet

Description

This command debugs ResvTear packets.

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

Parameters

detail

displays detailed information about ResvTear packets

srefresh

Syntax

srefresh [detail]

no srefresh

Context

debug>router>rsvp>packet

Description

This command debugs srefresh packets.

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

Parameters

detail

displays detailed information about RSVP-TE srefresh packets

4 PCEP

This section contains information about the following topics:

- [Introduction to the Path Computation Element Communication Protocol](#)
- [Base implementation of path computation elements \(PCE\)](#)
- [PCEP session establishment and maintenance](#)
- [PCEP parameters](#)
- [PCEP support for RSVP-TE LSPs](#)
- [LSP path diversity and bidirectionality constraints with path profiles](#)
- [NSP and VSR-NRC PCE redundancy](#)
- [Configuring RSVP-TE LSPs with PCEP using the CLI](#)
- [PCEP configuration command reference](#)

4.1 Introduction to the Path Computation Element Communication Protocol



Note: The 7705 SAR operates as a path computation client (PCC) only, supporting PCC capabilities for RSVP-TE LSPs and SR-TE LSPs. References to PCE router operation apply to the 7750 SR product family and are included for informational purposes only.

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 Nokia WAN SDN Controller is known as the Network Services Platform (NSP). The NSP is a set of applications which are 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) image 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 (NSP 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 illustrated in the following figure.

Figure 17: 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 providing 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 NRC-X.

The Network Resource Controller - Packet (NRC-P) implements the stateful path computation element (PCE) for packet networks. The following figure illustrates the NRC-P architecture and its main components.

Figure 18: 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 NSP 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 clients (PCCs). It also uses PCEP to communicate 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 segment routing (SR) or 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 per RFC 5440
- active and passive stateful PCE LSP update, as per *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 per *draft-ietf-pce-stateful-pce*
- support for LSP path diversity across different LERs using extensions to the PCE path profile, as per *draft-alvarez-pce-path-profiles*
- support for LSP path bidirectionality constraints using extensions to the PCE path profile, as per *draft-alvarez-pce-path-profiles*

4.2 Path computation fallback for PCC-initiated LSPs

For PCC-initiated SR-TE LSPs, the router supports fallback to a local path computation method (**fallback-path-computation-method local-cspf**) when the configured PCEP sessions are down or the PCE is unreachable. Fallback to a local computation method can be used when all configured PCEs are signaling overload and the redelegation timer has expired so that the LSP cannot be redelegated.

This capability is supported by both active stateful and passive stateful LSPs. Active stateful LSPs are fully delegated to the PCE by being both PCE-computed (**path-computation-method pce**) and PCE-controlled. Passive stateful LSPs are PCE-computed.



Note: For passive stateful LSPs, it is important that the **retry-timer** and **retry-limit** values exceed the **redelegation-timer** value, otherwise; the LSP may go operationally down before the fallback path computation has occurred.

A fallback path computation method is configured as follows:

```
configure>router>
mpls
  lsp lsp-name
    fallback-path-computation-method [local-cspf | none]
    path-computation-method pce
    pce-control
```

If the fallback path computation method is configured as **none**, MPLS uses the default method based on the configured path, which is hop-to-label path computation for SR-TE LSPs.

The **fallback-path-computation-method** command is only valid if **path-computation-method pce** is enabled.

The fallback mechanism is only triggered if the PCC informs MPLS that PCEP is down. It is not triggered while the PCC is administratively down or if it is not yet configured.



Note: On the first local path computation following a fallback, MPLS is not aware of the list of SRLGs or administrative groups that are used by the original path computed by PCE. As a result, MPLS can only provide a list of hops or links to avoid on the first computation.

PCE reports are sent, where applicable, with the delegation bit cleared.

4.3 Base implementation of path computation elements (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 32: 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 ¹
METRIC Object	N/A	PCReq, PCRep, PCRpt ¹
Explicit Route Object (ERO)	N/A	PCRep
Reported Route Object (RRO)	N/A	PCReq
LSPA Object	N/A	PCReq, PCRep, PCRpt ¹
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

Note:

1. Nokia proprietary

The behaviors and limitations of the implementation of the objects in the 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 will always set 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 will always set 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 will ignore them regardless of the setting of the P-flag, and will process 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 7705 SAR:

- PCE discovery using IS-IS, per RFC 5089, and OSPF, per RFC 5088, along with corresponding extensions for discovering stateful PCE, per *draft-sivabalan-pce-disco-stateful*
- security of the PCEP session using MD5 or TLS between PCEP peers
- PCEP synchronization optimization as per *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 per RFC 7471 and *draft-ietf-isis-te-metric-extensions*, and their use in the PCE metric object as per *draft-ietf-pce-pcep-service-aware*

4.4 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.


By default, the PCC checks first to determine if the remote PCE address is reachable out-of-band via the management port. If the remote address is not reachable, the PCC will try to reach the remote PCE address in-band. To restrict the PCC to only out-of-band or only in-band connection attempts, the **route-preference** command can be used. When the session comes up out-of-band, the system IP address is always used as the local address; 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 if 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 PCEP session initialization steps are illustrated in the following figure.

Figure 19: 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 re-established. Any subsequent change to the state of an LSP is synchronized at the time the session is re-established.

When a PCEP session to a peer times out or closes, the rate at which the PCEP speaker attempts the establishment of the session is subject to an exponential back-off mechanism.

4.5 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, meaning that each peer can 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, meaning that each PCEP speaker can propose a different dead timer value to its peer to use 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 re-establishment 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 re-establishing the PCEP session. When the PCEP session is re-established, 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.

4.5.1 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)

4.5.2 Stateful PCE

The main function introduced by stateful PCE over the base PCE implementation is the ability to synchronize the LSP state between the PCC and the PCE. This allows the PCE to have all the required LSP information to perform reoptimization and updating of the LSP paths.

The following table describes the TLVs, objects, and messages supported by stateful PCE on the 7705 SAR.

Table 33: PCEP stateful PCE extension TLVs, objects, and messages

TLV, object, or message	Contained in object	Contained in message
Path Computation State Report (PCRpt)	N/A	New message
Path Computation Update Request (PCUpd)	N/A	New message
STATEFUL-PCE-CAPABILITY TLV	OPEN	OPEN
Stateful Request Parameter (SRP) Object	N/A	PCRpt ¹ , PCErr

TLV, object, or message	Contained in object	Contained in message
LSP Object	ERO	PCRpt ¹ , PCReq, PCRep
LSP-IDENTIFIERS TLV	LSP	PCRpt ¹
SYMBOLIC-PATH-NAME TLV	LSP, SRP	PCRpt ¹
LSP-ERROR-CODE TLV	LSP	PCRpt ¹
RSVP-ERROR-SPEC TLV	LSP	PCRpt ¹

Note:**1. Nokia proprietary**

The behaviors and limitations of the implementation of the TLVs, objects, and messages in the table are as follows:

- The PCC and 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 specific PCEP session.
- The PCC always includes the LSP object in the PCReq message to make sure that the PCE can correlate the PLSP-ID for this LSP when a subsequent PCRpt message arrives with the delegation bit set. The PCE will, however, still honor a PCReq message without the LSP Object.
- PCE path computation will only consider the bandwidth used by LSPs in its LSP-DB. As a result, there are two situations where PCE path computation will not accurately take into account the bandwidth used in the network:
 - when there are LSPs that are signaled by the routers but are not synchronized with the PCE. The user can enable the reporting of the LSP to the PCE LSP database for each LSP.
 - when the stateful PCE is peering with a third-party stateless PCC, implementing only the original RFC 5440. While the PCE will be able to bring the PCEP session up, the LSP database will not be updated, because stateless PCC does not support the PCRpt message. Therefore, PCE path computation will not accurately take into account the bandwidth used by these LSPs in the network.
- The PCE ignores the R-flag (reoptimize flag) in the PCReq message when acting in passive stateful mode for an LSP and will always return the newly computed path, regardless whether it is link-by-link identical or has the same metric as the current path. The decision whether to initiate the new path in the network belongs to the PCC.
- The synchronization vector (SVEC) object is not supported on the 7705 SAR and the NSP. If the PCE receives a PCReq message with the SVEC object, it will ignore the SVEC object and treat each path computation request in the PCReq message as independent, regardless of the setting of the P-flag in the SVEC object common header.
- When an LSP is delegated to the PCE, there may be no prior state in the NRC-P LSP database for the LSP. This could be due to the PCE not having received a PCReq message for the same PLSP-ID. In order for the PCE to become aware of the original constraints of such an LSP, the following additional procedures are performed. These procedures are proprietary to Nokia:
 - The PCC appends a duplicate of each of the LSPA, metric, and bandwidth objects in the PCRpt message. The only difference between the two objects of the same type is that the P-flag is set in

the common header of the duplicate object to indicate 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 object and bandwidth object only. The 7705 SAR PCC does not support putting a bound value on the IGP or TE metric in the path computation.
- The path computation on the PCE uses the first set of objects when updating a path if the PCRpt message contains a single set. If the PCRpt message contains a duplicate set, PCE path computation must use the constraints in the duplicate set.
- For interoperability, implementations compliant with PCEP standards should be able to accept the first metric object and ignore the second object without additional error handling. Because there are also bandwidth and LSPA objects, the **[no] report-path-constraints** command is provided in the PCC on a per-PCEP session basis to disable the inclusion of the duplicate objects. Duplicate objects are included by default.

4.5.3 PCEP extensions in support of SR-TE LSPs

In order for the PCE and PCC to manage the path of an SR-TE LSP, they both implement the following extensions to PCEP in support of segment routing:

- The PCC and PCE support a Segment Routing capability TLV in the OPEN object to indicate support of segment routing tunnels by the PCE and the PCC during PCEP session initialization. This TLV is referred to as the SR-PCE-CAPABILITY TLV.
- The PCC and PCE support all PCEP capability TLVs defined in this chapter 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 the PCC will not use them during that specific PCEP session.
- The PCC and PCE support a new Path Setup Type TLV for SR-TE LSPs to be included in the Stateful PCE Request Parameters (SRP) object during path report (PCRpt) messages by the PCC.

A Path Setup Type TLV with a value of 1 identifies an SR-TE LSP.

- The PCC and PCE support a new Segment Routing ERO and RRO with sub-objects, referred to as SR-ERO and SR-RRO sub-objects, which encode the SID information in PCRpt messages.
- The PCE implementation supports the Segment-ID (SID) Depth value in the METRIC object. This is always signaled by the PCC in the PCEP Open object as part of the SR-PCE-CAPABILITY TLV. It is referred to as the Maximum Stack Depth (MSD). In addition, the per-LSP value for the **max-sr-labels** option, if configured, is signaled by the PCC to the PCE in the SID Depth value in a METRIC object for both a PCE-computed LSP and a PCE-controlled LSP. The PCE will compute and provide the full explicit path with TE links specified. If there is no path with the number of hops lower than the MSD value or the SID Depth value (if signaled), a reply with no path will be returned to the PCC. For a PCC-controlled LSP, if the label stack returned by the TE-DB hop-to-label translation exceeds the per-LSP maximum SR label stack size, the LSP is brought down.
- If the Path Setup Type (PST) TLV is not included in the PCReq message, the PCE or PCC must assume it is for an RSVP-TE LSP.

The following table describes the segment routing extension objects and TLVs supported on the 7705 SAR.

Table 34: PCEP segment routing extension objects and TLVs

TLV, object, or message	Contained in object	Contained in message
SR PCE CAPABILITY TLV	OPEN	OPEN
Path Setup Type (PST) TLV	SRP	PCReq, PCRep, PCRpt ¹
SR-ERO sub-object	ERO	PCRep, PCRpt ¹
SR-RRO sub-object	RRO	PCReq, PCRpt ¹
Segment-ID (SID) Depth Value in METRIC object	METRIC	PCReq, PCRpt ¹

Note:

1. Nokia proprietary

4.5.4 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 7705 SAR 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 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 **path-computation-method pce** 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.

4.5.5 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 the CLI or via the OSS/NSP NFM-P.
The configuration dictates which PCE control mode is required: active (**pce-control** and **pce-report** options enabled) or passive (**path-computation-method pce** enabled and **pce-control** disabled).
2. The 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. Note however that in the 7705 SAR, 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: For the PCE to be able 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 make sure that the **traffic-engineering** option is enabled in IGP. This causes IGP to flood the link SRLG and admin-group membership in its participating area, and for PCE to learn it in its TE database.

6. The PE router updates the CPM and the data path with the new path.
Up to this point, 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 to track error and state messages specific to this new path.
9. The PE router updates the CPM and the data path 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 data path, the synchronization of the PCC and PCE LSP databases, and the return of delegation to PCE.

The above 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 enabled the **path-computation-method pce** 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 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.

4.6 PCEP over TLS

PCEP over TLS (PCEPS) uses Transport Layer Security (TLS) to provide secure communication for PCEP. The secured session uses port 4189. The PCC is configured with a TLS client profile to initiate the TLS handshake. The PCE is configured with a TLS server profile to allow PCEP over TLS.

When a TLS server profile is configured on the PCE, the PCE can establish TLS connections in PCE secured (PCES) mode or non-TLS connections in PCE mode.

On the PCC, the 7705 SAR supports only strict TLS. That is, when the PCC is configured with a TLS client profile, the PCE must be TLS-capable and in PCES mode in order to establish a TLS connection. The PCE and PCC must perform a successful TLS handshake before the TLS wait timer expires.

If the PCC is not configured with a TLS client profile, the PCC and PCE can still make a connection in PCE mode, even if the PCE is configured with a TLS server profile.

In PCES mode, both the PCC and PCE must provide certificates for authentication. The PCE provides the server certificate to the PCC and requires the client certificate to authenticate the PCC.

4.6.1 TLS handshake

The 7705 SAR supports TLS client (PCC) functionality as well as TLS bidirectional authentication, where the PCE requests the client certificate to authenticate the PCC.

In a typical TLS handshake, the client starts the handshake with a ClientHello message. The server provides the server certificate for authentication to the client and sends a list of server-accepted ciphers.


The server can optionally ask the client to provide the client certificate using the server CertificateRequest option. When this option is present, the client provides the server with the client certificate, and if the certificate is authenticated, the TLS symmetric key is negotiated and the TLS session is established. The symmetric key is used to encrypt the TLS datapath.

See the *7705 SAR System Management Guide* for more information about the TLS handshake steps.

4.6.2 PCEP session over TLS

To establish a PCEP session over TLS as specified in RFC 8253, the PCC sends a StartTLS message to the PCE to initiate the TLS negotiation. The PCC activates the StartTLS timer and waits for the StartTLS message from the PCE. The timer is configured using the **config>router>pcep>pcc>peer>tls-wait-timer** command; the default timer is 60 s.

If the PCE is TLS-capable and sends back a StartTLS message before the StartTLS timer expires, the TLS handshake is initiated. If the PCE sends an Open message or does not send back a StartTLS message, the PCC responds with an error message, closes the TCP connection, and tries to establish the connection again. The PCEP Message-Type field of the PCEP common header for the StartTLS message is set to 13, as specified in RFC 8253. The following figure shows the establishment of a PCEP session over TLS.



TLS supports both in-band and out-of-band PCE connections.



Note: PCEP over TLS does not support CSM redundancy. After an activity switch, the PCEP over TLS connection goes down and must be reestablished.

4.7 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)

This section includes the following topics:

- [RSVP-TE LSP configuration for a PCC router](#)
- [Behavior of the LSP path update](#)
- [Behavior of LSP MBB](#)
- [Behavior of secondary LSP paths](#)
- [PCE path profile support](#)

4.7.1 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 path computation client (PCC). See [MPLS and RSVP-TE command reference](#) for command descriptions:

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

If the LSP has disabled PCE reporting, either due to inheritance from the MPLS-level configuration or due to 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 **one-hop-p2p** or a **mesh-p2p** RSVP-TE **auto-lsp** only supports the **pce-report** command in the LSP template:

```
config>router>mpls>lsp-template>pce-report {enable | disable | inherit}
```

The user must first shut down the LSP template before changing the value of the **pce-report** option.

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](#).

4.7.2 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 due to 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 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 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 data path, 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 an RSVP-TE LSP has the **no adaptive** option set, 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).

4.7.2.1 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 due to 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

4.7.3 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 config change MBB is only performed if the **path-computation-method pce**, **pce-report**, or **pce-control** options are enabled on the LSP. Otherwise, no action occurs. When **path-computation-method pce**, **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 and auto-bandwidth MBB types.

When the LSP goes into an MBB state at the ingress LER, the behavior is dependent on the operating mode of the LSP.

This section contains information about the following LSP MBB procedures:

- [PCC-controlled LSPs](#)
- [PCE-computed LSPs](#)
- [PCE-controlled LSPs](#)

4.7.3.1 PCC-controlled LSPs

All MBB types are supported for PCC-controlled LSPs. The LSP MBB procedures for a PCC-controlled LSP (**path-computation-method pce** 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.

4.7.3.2 PCE-computed LSPs

All MBB types are supported for PCE-computed LSPs. The LSP MBB procedures for a PCE-computed LSP (**path-computation-method pce** 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 **path-computation-method pce** to **no path-computation-method**, 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.

4.7.3.3 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, manual resignal, and auto-bandwidth 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 **path-computation-method pce** disabled, MPLS submits a path request to the local CSPF, which includes the updated path constraints.
3. For an LSP with **path-computation-method pce** 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 due to 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 due to 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, 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 functions as 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.

4.7.4 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 disjointness 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.

4.7.5 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 disjointness 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.

When the user wants to apply path disjointness 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 vice versa.

4.8 LSP path diversity and bidirectionality constraints with path profiles

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. As an alternative, the 7705 SAR supports the use of association groups to signal path diversity constraints. See [PCEP association groups](#) for information about association group support.

The configuration of an association group for an LSP and the configuration of the path profile ID and path group ID are mutually exclusive. A path profile and an association group are not interoperable. That is, LSPs with a path profile and those with an association group cannot be considered in the same group and cannot be compared against each other.

Path profiles are defined by the user 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 which PCE objective to use to optimize the path of the LSP in the PCE path profile, one of:

- IGP metric (cost)
- TE metric
- hops (span)

The CSPF algorithm optimizes 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 that the LSP belongs to. The CLI commands for this are described in [Configuring RSVP-TE LSPs with PCEP using the CLI](#). 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.

In order for the PCE to be 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 describes the new objects and TLVs introduced in the PCE path profile.

Table 35: 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 ¹

Note:

- 1. Nokia proprietary

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 chapter 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.

4.9 PCEP association groups

The 7705 SAR supports the use of PCEP association groups to reference SR-TE LSP path constraints. PCEP association groups allow LSPs to share common information, such as common policies or common configuration parameters. Users can use association groups for PCC-initiated SR-TE LSPs.

The PCEP ASSOCIATION object defines an Association ID and an Association Type to signal any type of association between LSPs (as defined in RFC 8697). Association groups are identified by a tuple consisting of an Association ID, Association Type, and Association Source. Association groups offer a disaggregated approach to specifying association, where the Association ID is equivalent to the path group ID, and the tuple (Association ID, Association Type, Association Source) is equivalent to the path profile ID. The tuple is the key to the association group. The IPv4 PCEP control plane supports signaling of Association IDs. The NSP uses the Association ID to reference a path profile.

Path diversity constraints are signaled using the Disjoint Association Type defined in RFC 8800, while policy constraints are signaled using the Policy Association Type defined in *draft-ietf-pce-association-policy-16*.

The Association ID and the Association Source address must be manually configured with the **diversity** or **policy** commands under the **config>router>pcep>pcc>pce-associations** context. In most cases, the Association Source address is the PCC source address. Configuring the Association Source address allows interoperability with third-party controllers that assume a specific allocation scheme, which supports cases such as the diversity between LSPs that originate on different PCC nodes but are configured with the same Association Source address for a specific Association ID and Association Type.

See [Diversity association groups](#) and [Policy association groups](#) for more details about these association types.

Association group support for PCC-initiated SR-TE LSPs

A PCC-initiated SR-TE can have up to five associations. These associations must be defined under the PCC context.

The configuration of an association group for an LSP and the configuration of the path profile ID and path group ID are mutually exclusive. A path profile and an association group are not interoperable. That is, LSPs with a path profile and those with an association group cannot be considered in the same group and cannot be compared against each other.

The PCC can signal one or more ASSOCIATION objects for each Association Type based on the configuration of the PCC-initiated LSP. This information is passed by MPLS to the PCC module. The PCC can include the ASSOCIATION object and its TLVs in PCEP PCReq and PCRpt messages. The PCE can reflect the same ASSOCIATION objects and TLVs in PCEP PCRep and PCUpd messages. For Diversity ASSOCIATION object, the PCE includes the status of the diversity as computed by the PCE in the DISJOINTNESS-STATUS-TLV.

ASSOCIATION object error handling

If the ASSOCIATION objects in the PCE do not match the ones sent by the PCC in the PCReq or PCRpt messages, the PCC returns an error.

If the PCC rejects a request for a particular association from MPLS with a "path not found" message, the PCC updates the reason code and tears down the path if it is up, then retries until the retry limit is exceeded.

The router ensures consistency between the association group configuration for a set of LSPs on a specific PCC. An association group can only have one set of parameters within an association (for example, Diversity Type and Disjointness Type), which ensures that LSPs added to the same association group do not have inconsistent parameters.

LSPs that originate on different PCCs can be added to the same association group even if the association group has different parameters configured on the different PCCs. In that case, only the NSP can detect parameter inconsistencies. For LSPs that are delegated and have inconsistent association parameters, the NSP sends a PCUpd down message followed by a PCErr message with the appropriate error message.

This causes the affected LSPs to go operationally down. For non-delegated LSPs, the NSP sends a PCErr message. If an association group with consistent parameters contains operationally up LSPs and a new LSP is added that has inconsistent parameters, the NSP only registers an error on the new LSP and leaves the existing LSPs undisturbed.

4.9.1 Diversity association groups

The DISJOINTNESS-CONFIGURATION TLV uses flag values to convey the diversity parameters requested for the set of SR-TE LSPs after the path computation. The same flags are set to indicate whether any of the parameters were met by the returned path.

Diversity association groups are configured with commands under the **config>router>pcep>pcc>pce-associations>diversity** context.

The **diversity-type** command is mandatory. If this command is not configured, the system does not validate the association configuration.

The **disjointness-type** command sets the T-flag in the DISJOINTNESS-CONFIGURATION TLV. This flag is used to determine the response if the PCE cannot provide a set of disjoint paths for one or more LSPs in the association.

The **disjointness-reference** command sets the P-flag in the DISJOINTNESS-CONFIGURATION TLV. This flag is used to indicate if this LSP path is the reference path for the disjoint set of paths. When set, the PCE must first compute the path of this LSP and then apply the configured disjointness type to compute the paths of all other LSPs in the same diversity association ID.

4.9.2 Policy association groups

Policy association groups are always operator-configured, as defined in *draft-ietf-pce-association-policy-16*. The value of the policy Association ID has a global meaning within a specific domain, independent of the source address in the ASSOCIATION object. For example, a PCC node can associate a policy Association ID with an SR-TE LSP included in the PCRpt message by including the PCC source address in the ASSOCIATION object, but the PCE looks up a path profile ID based solely on the Association ID value for PCC nodes in that domain. If another PCC node in the same domain signals another LSP with the same policy Association ID, the PCE also looks up the same path profile.

Policy association groups are configured with commands under the **config>router>pcep>pcc>pce-associations>policy** context.

4.10 NSP and VSR-NRC PCE redundancy

This feature introduces redundancy support to the PCE and PCC capabilities.

4.10.1 Feature configuration

A CLI command parameter has been introduced in the PCC for configuring the PCEP session to the standby backup PCE peer. A **preference** parameter value is used to specify the primary and the standby backup PCE peer in the **config>router>pcep>pcc** context:

```
peer ip-address preference preference
```


A maximum of two PCE peers are supported. The PCE peer that is not in overload is always selected by the PCC as the active PCE. However, if neither of the PCEs are signaling the overload state, the PCE with the higher numerical preference value is selected. In case of a tie, the PCE with the lower IP address is selected.

In order to change the value of the **preference** parameter, the peer must be deleted and recreated.

4.10.2 NSP ecosystem redundancy

The following figure illustrates the NSP ecosystem and the provision of redundancy across two separate sites. This is referred to as disaster recovery (DR) and is also sometimes referred to as geo-redundancy.

Figure 20: NSP ecosystem redundancy



NSP ecosystem redundancy consists of two mechanisms that can be deployed separately or together:

- high-availability (HA) at a single site:
 - the NSP, where the applications reside, is protected by a cluster of three VMs
 - the VSR-NRC module, which implements PCEP, OpenFlow, and BGP-LS/IGP, does not support HA and is deployed with a single VM that contains the combined CPM and IOM codes
- DR, which consists of a primary site and a secondary standby backup site. Each site consists of an NSP cluster and a VSR-NRC VM complex. A heartbeat protocol runs between the NSP clusters at the primary site and the standby backup site.

The VSR-NRC can be deployed as a standalone configuration; however, the NSP must be deployed in a cluster at each site. This configuration is also referred to as a 3+3 deployment.

Each parent NSP cluster establishes a reliable TCP session with a virtual IP to the local VSR-NRC. The TCP session runs an internal protocol, also known as cproto. This configuration is done prior to system startup and cannot be changed with an active NSP; the NSP must be shut down for any changes.

4.10.2.1 NSP cluster redundancy

The following describes NSP cluster redundancy rules:

- At a single site, a master is elected among the cluster of three VMs. Between sites, a single cluster at one site is the primary/active site and the other DR site is the secondary/standby site.
- The application processes at the standby site are shut down but the Neo4j and other databases are synchronized with the primary/active site.
- Switching to the standby site can be initiated manually or by using an automated approach stemming from the loss of heartbeat between the primary and standby sites.
- If the NSP cluster at the primary/active site is down (two out of three servers must be inactive, shut down, or failed), the heartbeat mechanism between the primary and standby NSP clusters fails after three timeouts. This initiates the active role at the secondary/standby site.
- When the NSP cluster at the primary site is back up, the heartbeat mechanism between the primary/standby and secondary/active NSP clusters is restored. The primary site can be manually restored as the active site. Automatic reversion to the primary NSP cluster is not supported.

4.10.2.2 VSR-NRC redundancy

The following describes VSR-NRC redundancy rules:

- **steady state behavior**
 - The VSR-NRC at the secondary/standby site, in the same way as the VSR-NRC at the primary/active site, establishes PCEP sessions to the PCCs. However, the VSR-NRC at the standby site has its PCEP sessions to the PCCs in the overload state. The VSR-NRC enters this PCEP overload state when its upstream cproto session to the NSP cluster is down, resulting from either the NSP cluster going into the standby state or the cproto session failing.
 - The VSR-NRC acting as a PCE signals the overload state to the PCCs in a PCEP notification message. In the overload state, the VSR-NRC PCE accepts reports (PCRpt) without delegation but rejects requests (PCReq) and rejects reports (PCRpt) with delegation. The VSR-NRC PCE also does not originate initiate messages (PCInitiate) or update messages (PCUpd).
 - The VSR-NRC at the secondary/standby site maintains its BGP and IGP peerings with the network and updates its TE database as a result of any network topology changes.
- **primary/active NSP cluster failure**

If the NSP cluster at the primary/active site is down (two out of three servers must be inactive, shut down, or failed), the heartbeat mechanism between the primary/active and secondary/standby NSP clusters fails. This initiates the NSP cluster activity at the secondary/standby site.

The following are the activities on the VSR-NRC:

- The VSR-NRC at the primary site detects the cproto session failure and puts all its PCEP sessions to the PCCs into the overload state.
- The NSP cluster at the secondary site establishes the cproto session to the local VSR-NRC, which then brings its PCEP sessions out of the overload state.
- The VSR-NRC at the secondary site begins synchronizing the TE and LSP databases with the parent NSP cluster at the secondary site that is now the active site.
- The VSR-NRC at the primary site must also return the delegation of all LSPs to the PCCs by sending an empty LSP Update Request that has the Delegate flag set to 0 as per RFC 8231. This allows the PCCs to delegate all eligible LSPs, including PCE-initiated LSPs, to the PCE function in the VSR-NRC at the secondary site.



Note: If the entire primary site fails, the above actions of the VSR-NRC at the primary site do not apply; however, the remaining actions do apply.

- **VSR-NRC complex failure at the primary site (NSP server is still up)**

A VSR-NRC complex failure at the primary/active NSP site does not initiate an NSP switchover to the secondary/standby NSP site. If the VSR-NRC at the primary site does not recover, a manual switchover to the secondary NSP site is required. The VSR-NRC failure causes a cproto session failure alarm to be raised on the NSP indicating that the NSP cannot communicate with the VSR-NRC. The operator can manually perform a switchover of the NSP activity to the secondary site.

4.10.2.3 PCC behavior

The following describes rules for a PCC with PCE redundancy:

- PCCs can establish upstream PCEP sessions with at most two VSR-NRC PCEs.
- Each upstream session has a preference that takes effect when both upstream PCEP sessions are successfully established. The PCE peer that is not in overload is always selected by the PCC as the active PCE. However, if neither of the PCEs are signaling the overload state, the PCE with the higher numerical preference value is selected, and in case of a tie, the PCE with the lower IP address is selected.
- In the steady state, because one upstream VSR-NRC PCE is in overload, only one PCEP session is active. The PCCs delegate an LSP using a report message (PCRpt) with the Delegate flag set to the active VSR-NRC PCE only. Request messages (PCReq) are not sent to the secondary/standby VSR-NRC PCE in overload. PCRpt messages are sent with the Delegate flag clear to the secondary/standby VSR-NRC PCE in overload.
- If the current active PCEP session signals the overload state, the PCC will select the other PCE as the active PCE as long as the corresponding PCEP session is not in overload. Any new path request message (PCReq) or path report message (PCRpt) with the Delegate flag set is sent to the new PCE. The PCE in overload should return the delegation of all existing LSPs to this PCC by sending an empty LSP Update Request that has the Delegate flag set as per RFC 8231. This PCC will then delegate these LSPs to the newly active PCE by sending a path report (PCRpt) with the Delegate flag set.
- If the current active PCEP session goes operationally down, the PCC starts the redelegation timer (default is 90 s).
 - If the PCEP session is restored before the redelegation timer expires, no delegation change is performed and the LSP state is maintained.
 - Upon expiration of the redelegation timer, the PCC looks for the other PCEP session and, if not in overload, it immediately delegates the LSPs to the newly active PCE. If the new PCE accepts the delegation, the LSP state is maintained.
 - If the PCEP session does not recover before the redelegation timer expires and the PCC fails to find another active PCEP session, the state of LSPs by default are not cleared. The PCC does not clear the state of PCC-initiated LSPs; however, the user can do this by deleting the configuration.

4.11 Configuring RSVP-TE LSPs with PCEP using the CLI

This section provides information about using the CLI to configure and operate RSVP-TE LSPs with PCEP.

The following information describes the detailed configuration of an inter-area RSVP-TE LSP with both a primary path and a secondary path. The network uses IS-IS with the backbone area in Level 2 and the leaf areas in Level 1. Topology discovery is learned by the NRC-P using IGP and the NSP NFM-P.

The LSP uses an admin-group constraint to keep the paths of the secondary and primary links disjoint in the backbone area. The LSP is PCE-controlled but also has **path-computation-method pce** enabled so that the initial path, and any MBB path, is also computed by the PCE.

The NSP and 7705 SAR load versions used to produce this example in this section are:

- NSP: NSP-2.0.3-rel.108
- PCE SROS: TiMOS-B-0.0.W129
- PCC: TiMOS-B-0.0.I4902


This section provides configuration and show commands for the following examples:

- [PCEP on the PCE node and the PCC node](#)
- [MPLS on the PCC node](#)

4.11.1 PCEP on the PCE node and the PCC node

The following figure shows a multi-level IS-IS topology in the NSP GUI.

Figure 21: Multi-level IS-IS topology in the NSP GUI



The following example shows the configuration and show command output of PCEP on the PCE node and the PCC node.



Note: The 7705 SAR operates as a path computation client (PCC) only, supporting PCC capabilities for RSVP-TE LSPs and SR-TE LSPs. References to PCE router operation apply to the 7750 SR product family and are included for informational purposes only.

```
*A:PCE Server 226>config>router>pcep>pce# info
-----
      local-address 192.168.48.226
      no shutdown
-----

*A:Reno 194>config>router>pcep>pcc# info
-----
      peer 192.168.48.226
      no shutdown
      exit
      no shutdown
-----

*A:PCE Server 226>config>router>pcep>pce# show router pcep pce status
=====
Path Computation Element Protocol (PCEP) Path Computation Element (PCE) Info
=====
Admin Status           : Up           Oper Status           : Up
Unknown Msg Limit      : 10 msg/min
Keepalive Interval     : 30 seconds   DeadTimer Interval    : 120 seconds
Capabilities List       : stateful-delegate stateful-pce segment-rt-path
Local Address          : 192.168.48.226
PCE Overloaded         : false
-----

PCEP Path Computation Element (PCE) Peer Info
-----
Peer                   Sync State           Oper Keepalive/Oper DeadTimer
-----
192.168.48.190:4189    done                 30/120
192.168.48.194:4189    done                 30/120
192.168.48.198:4189    done                 30/120
192.168.48.199:4189    done                 30/120
192.168.48.219:4189    done                 30/120
192.168.48.221:4189    done                 30/120
192.168.48.224:4189    done                 30/120
-----
=====

*A:Reno 194# show router pcep pcc status
=====
Path Computation Element Protocol (PCEP) Path Computation Client (PCC) Info
=====
Admin Status           : Up           Oper Status           : Up
Unknown Msg Limit      : 10 msg/min
Keepalive Interval     : 30 seconds   DeadTimer Interval    : 120 seconds
Capabilities List       : stateful-delegate stateful-pce segment-rt-path
Address                : 192.168.48.194
Report Path Constraints: True
-----

PCEP Path Computation Client (PCC) Peer Info
-----
Peer                   Admin State/Oper State Oper Keepalive/Oper DeadTimer
-----
192.168.48.226         Up/Up                 30/120
-----
=====

*A:Reno 194# show router pcep pcc lsp-db
=====
```

```


PCEP Path Computation Client (PCC) LSP Update Info
=====
PCEP-specific LSP ID: 11
LSP ID       : 14378                LSP Type       : rsvp-p2p
Tunnel ID    : 1                    Extended Tunnel Id : 192.168.48.194
LSP Name     : From Reno to Atlanta RSVP-TE::primary_empty
Source Address : 192.168.48.194    Destination Address : 192.168.48.224
LSP Delegated : True              Delegate PCE Address: 192.168.48.226
Oper Status  : active
-----
PCEP-specific LSP ID: 12
LSP ID       : 14380                LSP Type       : rsvp-p2p
Tunnel ID    : 1                    Extended Tunnel Id : 192.168.48.194
LSP Name     : From Reno to Atlanta RSVP-TE::secondary_empty
Source Address : 192.168.48.194    Destination Address : 192.168.48.224
LSP Delegated : True              Delegate PCE Address: 192.168.48.226
Oper Status  : up
=====

```

4.11.2 MPLS on the PCC node

The following figure shows primary and secondary RSVP-TE LSP paths in the NSP GUI.

Figure 22: Primary and secondary RSVP-TE LSP paths in the NSP GUI



The following example shows the configuration and show command output of MPLS on the PCC node.

```

*A:Reno 194>config>router>mpls>lsp# info
-----
to 192.168.48.224
egress-statistics
shutdown
exit
fast-reroute facility
no node-protect
exit
path-computation-method pce
pce-report enable
pce-control
revert-timer 1
primary "primary_empty"

```



```

        exclude "top"
        bandwidth 10
    exit
    secondary "secondary_empty"
        standby
        exclude "bottom"
        bandwidth 5
    exit
    no shutdown
-----

*A:Reno 194# show router mpls lsp "From Reno to Atlanta RSVP-TE" path detail
=====
MPLS LSP From Reno to Atlanta RSVP-TE Path (Detail)
=====
Legend :
  @ - Detour Available          # - Detour In Use
  b - Bandwidth Protected       n - Node Protected
  s - Soft Preemption
  S - Strict                    L - Loose
  A - ABR
=====
-----
LSP From Reno to Atlanta RSVP-TE Path primary_empty
-----
LSP Name       : From Reno to Atlanta RSVP-TE
Path LSP ID    : 14382
From           : 192.168.48.194      To           : 192.168.48.224
Admin State    : Up                  Oper State    : Up
Path Name      : primary_empty       Path Type     : Primary
Path Admin     : Up                  Path Oper     : Up
Out Interface  : 1/1/1               Out Label     : 262094
Path Up Time   : 0d 00:00:22         Path Down Time : 0d 00:00:00
Retry Limit    : 0                   Retry Timer    : 30 sec
Retry Attempt  : 0                   Next Retry In  : 0 sec
BFDD Template  : None                BFDD Ping Interval : 60
BFDD Enable    : False
Adspec         : Disabled            Oper Adspec    : Disabled
CSPF           : Enabled             Oper CSPF       : Enabled
Least Fill     : Disabled            Oper LeastFill  : Disabled
FRR            : Enabled             Oper FRR        : Enabled
FRR NodeProtect : Disabled           Oper FRR NP     : Disabled
FR Hop Limit   : 16                  Oper FRHopLimit : 16
FR Prop Admin Gr*: Disabled          Oper FRPropAdmGrp : Disabled
Propagate Adm Grp: Disabled          Oper Prop Adm Grp : Disabled
Inter-area     : False
PCE Updt ID    : 0
PCE Report     : Enabled             Oper PCE Report : Enabled
PCE Control    : Enabled             Oper PCE Control : Enabled
PCE Compute    : Enabled
Neg MTU        : 1496                Oper MTU        : 1496
Bandwidth      : 10 Mbps              Oper Bandwidth   : 10 Mbps
Hop Limit      : 255                  Oper HopLimit    : 255
Record Route   : Record               Oper Record Route : Record
Record Label   : Record               Oper Record Label : Record
Setup Priority  : 7                    Oper Setup Priority : 7
Hold Priority   : 0                    Oper Hold Priority : 0
Class Type     : 0                    Oper CT          : 0
Backup CT      : None
MainCT Retry   : n/a
Rem            :
MainCT Retry   : 0
Limit          :
Include Groups :                      Oper Include Groups :

```

```

None
Exclude Groups :
top
Adaptive : Enabled
Preference : n/a
Path Trans : 7
Failure Code : noError
Failure Node : n/a
Explicit Hops :
  No Hops Specified
Actual Hops :
  10.202.5.194 (192.168.48.194) @
  -> 10.202.5.199 (192.168.48.199) @
  -> 192.168.48.185 (192.168.48.185)
  -> 10.0.5.185
  -> 192.168.48.223 (192.168.48.223)
  -> 10.0.7.223
  -> 192.168.48.224 (192.168.48.224)
  -> 10.101.4.224
Computed Hops :
  10.202.5.199(S)
  -> 10.0.5.185(S)
  -> 10.0.7.223(S)
  -> 10.101.4.224(S)
Resignal Eligible: False
Last Resignal : n/a
CSPF Metric : 40
-----
LSP From Reno to Atlanta RSVP-TE Path secondary_empty
-----
LSP Name : From Reno to Atlanta RSVP-TE
Path LSP ID : 14384
From : 192.168.48.194 To : 192.168.48.224
Admin State : Up Oper State : Up
Path Name : secondary_empty Path Type : Standby
Path Admin : Up Path Oper : Up
Out Interface : 1/1/1 Out Label : 262091
Path Up Time : 0d 00:00:25 Path Down Time : 0d 00:00:00
Retry Limit : 0 Retry Timer : 30 sec
Retry Attempt : 0 Next Retry In : 0 sec
BFD Template : None BFD Ping Interval : 60
BFD Enable : False
Adspec : Disabled Oper Adspec : Disabled
CSPF : Enabled Oper CSPF : Enabled
Least Fill : Disabled Oper LeastFill : Disabled
Propagate Adm Grp: Disabled Oper Prop Adm Grp : Disabled
Inter-area : False
PCE Updt ID : 0
PCE Report : Enabled Oper PCE Report : Enabled
PCE Control : Enabled Oper PCE Control : Enabled
PCE Compute : Enabled
Neg MTU : 1496 Oper MTU : 1496
Bandwidth : 5 Mbps Oper Bandwidth : 5 Mbps
Hop Limit : 255 Oper HopLimit : 255
Record Route : Record Oper Record Route : Record
Record Label : Record Oper Record Label : Record
Setup Priority : 7 Oper Setup Priority : 7
Hold Priority : 0 Oper Hold Priority : 0
Class Type : 0 Oper CT : 0
Include Groups :
None
Exclude Groups :
bottom
Adaptive : Enabled Oper Include Groups :
Preference : 255 Oper Exclude Groups :
bottom
Oper Metric : 60

```

```

Path Trans      : 28                      CSPF Queries      : 10
Failure Code    : noError
Failure Node    : n/a
Explicit Hops   :
  No Hops Specified
Actual Hops     :
  10.202.5.194 (192.168.48.194)           Record Label      : N/A
-> 10.202.5.199 (192.168.48.199)           Record Label      : 262091
-> 10.0.9.198 (192.168.48.198)            Record Label      : 262096
-> 192.168.48.184 (192.168.48.184)         Record Label      : 262102
-> 10.0.2.184                             Record Label      : 262102
-> 192.168.48.221 (192.168.48.221)         Record Label      : 262119
-> 10.0.4.221                             Record Label      : 262119
-> 192.168.48.223 (192.168.48.223)         Record Label      : 262088
-> 10.0.10.223                            Record Label      : 262088
-> 192.168.48.224 (192.168.48.224)         Record Label      : 262115
-> 10.101.4.224                           Record Label      : 262115
Computed Hops   :
  10.202.5.199(S)
-> 10.0.9.198(S)
-> 10.0.2.184(S)
-> 10.0.4.221(S)
-> 10.0.10.223(S)
-> 10.101.4.224(S)
Srlg            : Disabled
Srlg Disjoint   : False
Resignal Eligible: False
Last Resignal   : n/a                      CSPF Metric        : 60
=====

```

For more information and CLI output examples, see the **show router mpls sr-te-lsp** command description.

4.12 PCEP configuration command reference

4.12.1 Command hierarchies

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

4.12.1.1 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
- pce-associations
- [no] diversity association-name
- association-id association-id
- no association-id
- association-source ip-address
- no association-source
- [no] disjointness-reference
- disjointness-type {loose | strict}
- no disjointness-type
- diversity-type {link | node | srlg-link | srlg-node}
- no diversity-type
- [no] policy association-name
- association-id association-id
- no association-id
- association-source ip-address
- no association-source
- peer ip-address [preference preference]
- no peer ip-address
- route-preference {both | inband | outband}
- [no] route-preference
- [no] shutdown
- tls-client-profile profile-name
- no tls-client-profile
- tls-wait-timer [60..255]
- no tls-wait-timer
- redelegation-timer seconds
- no redelegation-timer
- [no] report-path-constraints
- [no] shutdown
- unknown-message-rate msg/min
- no unknown-message-rate

```

4.12.1.2 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]
- pce-associations
-   diversity [diversity-assoc-name]
-   policy [policy-assoc-name]
- peer [ip-address] [detail]
- status
```

4.12.2 Command descriptions

- [PCEP commands](#)
- [Show commands](#)

4.12.2.1 PCEP commands

pcep

Syntax

[no] pcep

Context

config>router

Description

This command enables the Path Computation Element Communication Protocol (PCEP) and enters the context to configure PCEP parameters.

The **no** form of the command disables PCEP.

pcc

Syntax

[no] pcc

Context

config>router>pcep

Description

This command enables the context to configure PCC parameters.

dead-timer

Syntax

dead-timer *seconds*

no dead-timer

Context

config>router>pcep>pcc

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, meaning that each PCEP speaker can propose a different dead timer value to its peer to use to detect session timeout.

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

Default

120

Parameters

seconds

the dead timer value, in seconds

Values 1 to 255

keepalive

Syntax

keepalive *seconds*

no keepalive

Context

config>router>pcep>pcc

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, meaning that each peer can use a different keepalive timer value at its end.

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

Default

30

Parameters

seconds

the keepalive value, in seconds

Values 1 to 255

local-address

Syntax

local-address *ip-address*

no local-address

Context

config>router>pcep>pcc

Description

This command configures the local address of the PCEP speaker.

The PCEP protocol 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 the connection is established, the PCC and PCE exchange OPEN messages, which initializes the PCEP session and exchanges the session parameters to be negotiated.

By default, the PCC checks first to determine if the remote PCE address is reachable out-of-band via the management port. If the remote address is not reachable, the PCC will check if the remote PCE address is reachable in-band. The local address configured by the user is only used for in-band sessions and is otherwise ignored. When the session comes up out-of-band, the system IP address is always used as the local address. Use the **route-preference** command to configure the PCC to use only out-of-band or only in-band instead.

The **no** form of the command removes the configured local address of the PCEP speaker.

Parameters

ip-address

the IP address of the PCEP speaker to be used for in-band sessions

pce-associations

Syntax

pce-associations

Context

config>router>pcep>pcc

Description

This command enables the context to configure PCEP association groups.

Default

n/a

diversity

Syntax

[no] diversity *association-name*

Context

config>router>pcep>pcc>pce-assoc

Description

This command creates a diversity association, and enables the context to create diversity association parameters .

The **no** form of the command deletes the diversity association.

Default

n/a

Parameters

association-name

the name of the diversity association, up to 32 characters

association-id

Syntax

association-id *association-id*

no association-id

Context

config>router>pcep>pcc>pce-assoc>div

config>router>pcep>pcc>pce-assoc>plcy

Description

This command configures the diversity or policy Association ID. The user must specify an Association ID. The ID value is interpreted as a global value independent of the IP address of the PCC node that uses it to associate LSPs.

The Association ID for the diversity type supports the user-configured mode of operation, as defined in RFC 8800.

The **no** form of the command removes the Association ID from the diversity or policy association.

Default

no association-id

Parameters

association-id

specifies the diversity or policy Association ID

Values 1 to 65535

association-source

Syntax

association-source *ip-address*

no association-source

Context

config>router>pcep>pcc>pce-assoc>div

config>router>pcep>pcc>pce-assoc>plcy

Description

This command configures the source IP address of the diversity or policy association.

The **no** form of the command removes the IP address from the diversity or policy association.

Default

no association-source

Parameters

ip-address

the source IP address

disjointness-reference

Syntax

[no] disjointness-reference

Context

config>router>pcep>pcc>pce-assoc>div

Description

This command sets the P-flag in the DISJOINTNESS-CONFIGURATION TLV. When enabled, it indicates that this LSP path is the reference path for the disjoint set of paths. The PCE must first compute the path of this LSP and then apply the configured disjointness type to compute the paths of all other LSPs in the same diversity Association ID.

The **no** form of the command sets the P-flag to false.

Default

no disjointness-reference

disjointness-type**Syntax**

disjointness-type {loose | strict}

no disjointness-type

Context

config>router>pcep>pcc>pce-assoc>div

Description

This command sets the T-flag in the DISJOINTNESS-CONFIGURATION TLV.

When the disjointness type is configured as **strict**, the T-flag is set in the DISJOINTNESS-CONFIGURATION TLV and the PCE will send a reply with no path if it is not able to provide a set of disjoint paths for one or more LSPs in the association.

When the disjointness type is configured as **loose**, the T-flag is not set in the DISJOINTNESS-CONFIGURATION TLV and the PCE is allowed to relax the disjointness constraint in order to return paths for the LSPs.

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

Default

loose

Parameters

loose

specifies the loose disjointness type

strict

specifies the strict disjointness type

diversity-type**Syntax**

diversity-type {link | node | srlg-link | srlg-node}

no diversity-type

Context

config>router>pcep>pcc>pce-assoc>div

Description

This command configures the diversity type for the association group. This command is mandatory. If the command is not configured, the system does not validate the association configuration.

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

Default

no diversity-type

Parameters

link

specifies the link diversity type

node

specifies the node diversity type

srlg-link

specifies the SRLG link diversity type

srlg-node

specifies the SRLG node diversity type

policy

Syntax

[no] policy *association-name*

Context

config>router>pcep>pcc>pce-assoc

Description

This command creates a PCE policy association and enables the context to create policy association parameters.

The **no** form of the command deletes the policy association.

Default

n/a

Parameters

association-name

the name of the policy association, up to 32 characters

peer

Syntax

peer *ip-address* [**preference** *preference*]

no peer *ip-address*

Context

config>router>pcep>pcc

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 **preference** parameter allows the PCC to select the preferred PCE when both have PCEP sessions successfully established. A maximum of two PCEP peers is supported. The PCE peer that is not in overload is always selected by the PCC as the active PCE; however, if neither of the PCEs are signaling the overload state, the PCE with the lower IP address is selected.

The **no** form of the command removes the specified peer PCEP speaker.

Parameters

ip-address

the IPv4 address of the PCEP peer to be used as the destination address in the PCEP session

preference

specifies the preference value of the peer

Values 0 to 100

Default 0

redelegation-timer

Syntax

redelegation-timer *seconds*

no redelegation-timer

Context

config>router>pcep>pcc

Description

This command configures the redelegation timer for PCE-initiated LSPs.

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

Default

90

Parameters*seconds*

specifies the number of seconds before the redelegation timer expires

Values 1 to 3600**route-preference****Syntax****route-preference** {**both** | **inband** | **outband**}**no route-preference****Context**

config>router>pcep>pcc>peer

Description

This command configures the route preference that PCC uses to reach the remote PCE address, which allows you to restrict the PCC to using only out-of-band or only in-band for reachability. By default, the PCC checks first to determine if the remote PCE address is reachable out-of-band via the management port. If the remote address is not reachable, the PCC will try to reach the remote PCE address in-band.

The **no** form of this command resets the route preference to the default behavior.

Default

both

Parameters**both**

the PCC uses the out-of-band routing table to reach the PCC; if unsuccessful, it then uses the in-band routing table

inband

the PCC uses only the in-band routing table to attempt to reach the PCE

outband

the PCC uses only the out-of-band management router routing table to attempt to reach the PCE

report-path-constraints**Syntax****[no] report-path-constraints**

Context

```
config>router>pcep>pcc
```

Description

This command enables the inclusion of LSP path constraints in the PCE report messages sent from the PCC to a PCE.

In order 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 7705 SAR 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 the command disables the above 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
```

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**

Default

shutdown

tls-client-profile

Syntax

tls-client-profile *profile-name*

no tls-client-profile

Context

config>router>pcep>pcc>peer

Description

This command configures a TLS client profile on the PCC. When the TLS client profile is configured, the PCC tries to establish a PCEP connection with the PCE over TLS. If the PCE sends an Open message or does not send back a StartTLS message within the time set by the **tls-wait-timer** command, the connection is not accepted and the TCP connection is closed and a new connection is tried within 60 s

The **no** form of the command removes TLS encryption from the communication between the PCC and the PCE.

Default

no tls-client-profile

Parameters

profile-name

the TLS client profile name, up to 32 characters in length

tls-wait-timer

Syntax

tls-wait-timer [*60..255*]

no tls-wait-timer

Context

config>router>pcep>pcc>peer

Description

This command configures the time that the PCC waits before declaring a TLS handshake failure if the handshake is not established.

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

Default

tls-wait-timer 60

Parameters

[60..255]

the wait time in seconds

Values: 60 to 255

unknown-message-rate

Syntax

unknown-message-rate *msg/min*

no unknown-message-rate

Context

config>router>pcep>pcc

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 the command returns the unknown message rate to the default value.

Default

10

Parameters

msg/min

the rate of unknown messages, in messages per minute

Values 1 to 255

4.12.2.2 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

Description
This command displays PCEP PCC detailed information.

Output
The following output is an example of PCEP PCC detailed information, and [Table 36: PCEP PCC field descriptions](#) describes the fields.

Output example

```
*A:Sar18 Dut-B>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:Sar18 Dut-B>show>router>pcep#
```

Table 36: PCEP PCC field descriptions

Label	Description
Admin Status	The administrative status of the PCC
Oper Status	The operational status of the PCC
Unknown Msg Limit	The maximum rate of unknown messages that can be received on a PCEP session

Label	Description
Keepalive Interval	The specified keepalive interval for the PCEP session
DeadTimer Interval	The specified dead time interval for the PCEP session
Capabilities List	The capabilities list for the PCEP session
Address	The local IP address of the PCEP speaker
Report Path Constraints	Indicates whether to include LSP path constraints in the PCE report messages sent from the PCC to a PCE
Open Wait Timer	The value of the open wait timer for the PCEP session
Keep Wait Timer	The value of the keep wait timer for the PCEP session
Sync Timer	The value of the synchronization timer for the PCEP session
Request Timer	The value of the request timer for the PCEP session
Connection Timer	The value of the keep wait timer for the PCEP session
Allow Negotiations	Indicates where negotiations between PCEP PCC and PCE are allowed
Max Sessions	The maximum number of PCEP sessions on the router
Max Unknown Req	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
```

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 0 to 65535

ip-address

specifies an IPv4 address

Output

The following output is an example of PCEP PCC LSP information.

Output example

```
*A:7705:Dut-C# 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: 192.120.210.36
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: 192.120.210.36
Oper Status      : up
=====
*A:7705:Dut-C#
```

path-request

Syntax

path-request [*lsp-type* {rsvp-p2p}] [*dest ip-address*] [*detail*]

Context

show>router>pcep>pcc

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
- detail

displays detailed path request information

Output

The following output is an example of PCEP PCC path request information.

Output example

```
*A:7705:Dut-C# 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
=====
```

diversity

Syntax

diversity [*diversity-assoc-name*]

Context

show>router>pcep>pcc>pce-assoc

Description

This command displays the diversity association information for a PCEP diversity association group.
If a diversity association name is not specified, the information for all configured diversity associations is displayed.

Parameters

- diversity-assoc-name*

the diversity association name

Output

The following output is an example of PCEP diversity association information, and [Table 37: PCEP diversity association field descriptions](#) describes the output fields.

Output example

```
*A:Dut-C>config>router>mpls>lsp# show router pcep pcc pce-associations diversity "test_div"
=====
PCEP PCC Diversity Associations Info
=====
Association Name : test_div
Association ID : 20
Diversity Type : link
Disjoint Reference : False
Association Source : 192.168.213.59
Disjointness Type : loose
=====
```

Table 37: PCEP diversity association field descriptions

Label	Description
Association Name	The diversity association name
Association ID	The diversity association ID
Association Source	The source IP address of the diversity association
Diversity Type	The diversity type for the diversity association: association: link, node, srlg-link, or srlg-node
Disjointness Type	The disjointness type for the diversity association: loose or strict
Disjoint Reference	Indicates whether the P-flag is set in the DISJOINTNESS CONFIGURATION TLV; false means that the flag is not set

policy

Syntax

policy [*policy-assoc-name*]

Context

show>router>pcep>pcc>pce-assoc

Description

This command displays the policy association information for a PCEP policy association group.

If a policy association name is not specified, the information for all configured policy associations is displayed.

Parameters

policy-assoc-name
the policy association name

Output

The following output is an example of PCEP policy association information, and [Table 38: PCEP policy association field descriptions](#) describes the output fields.

Output example

```
*A:Dut-C>config>router>mpls>lsp# show router pcep pcc pce-associations policy "test_pol"
=====
PCEP PCC Policy Associations Info
=====
Association Name : test_pol
Association ID : 10                      Association Source : 192.168.213.59
=====
```

Table 38: PCEP policy association field descriptions

Label	Description
Association Name	The policy association name
Association ID	The policy association ID
Association Source	The source IP address of the policy association

peer

Syntax

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

Context

show>router>pcep>pcc

Description

This command displays PCEP PCC peer information.

Parameters

- ip-address*
specifies a peer IPv4 address to display
- detail**
displays detailed peer information

Output

The following output is an example of a PCEP PCC peer information, and [Table 39: PCEP PCC peer field descriptions](#) describes the fields.

Output example

```
*A:Sar18 Dut-B>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: Sar18 Dut-B>show>router>pcep>pcc#

```

Table 39: PCEP PCC peer field descriptions

Label	Description
IP Address	The IP address of the PCC peer
Admin Status	The administrative status of the PCC peer
Oper Status	The operational status of the PCC peer
Peer Capabilities	The PCEP capabilities of the PCC peer
Speaker ID	The IP address of the PCC peer speaker
Sync State	The synchronization state of the
Peer Overloaded	Indicates whether the PCC peer is overloaded
Session Establish Time	The length of time since the PCEP session was established
Oper Keepalive	The operational value for the PCC peer keepalive timer
Oper DeadTimer	The operational value for the PCC peer dead timer
Session Setup Count	The number of times that the PCEP session has been set up
Session Setup Fail Count	The number of times that the PCEP session failed to be set up

Label	Description
Statistics Information	
PC Request Message	The number of path computation (PC) request messages sent the PCC peer and received from the PCC peer
PC Reply Message	The number of PC reply messages sent to the PCC peer and received from the PCC peer
PC Error Message	The number of PC error messages sent to the PCC peer and received from the PCC peer
PC Notification Message	The number of PC notification messages sent to the PCC peer and received from the PCC peer
PC Keepalive Message	The number of PC keepalive messages sent to the PCC peer and received from the PCC peer
PC Update Message	The number of PC update messages sent to the PCC peer and received from the PCC peer
PC Report Message	The number of PC report messages sent to the PCC peer and received from the PCC peer
Path Report	The number of path reports sent to the PCC peer and received from the PCC peer
Path Request	The path requests sent to the PCC peer and received from the PCC peer

status

Syntax
status

Context
show>router>pcep>pcc

Description
This command displays PCEP PCC status information.

Output
The following output is an example of a PCEP PCC status information, and [Table 40: PCEP PCC status field descriptions](#) describes the fields.

Output example

```
*A: Sar18 Dut-B>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: Sar18 Dut-B>show>router>pcep>pcc#

```


Table 40: PCEP PCC status field descriptions

Label	Description
Admin Status	The administrative status of the PCC
Oper Status	The operational status of the PCC
Unknown Msg Limit	The maximum rate of unknown messages that can be received on a PCEP session
Keepalive Interval	The specified keepalive interval for the PCEP session
DeadTimer Interval	The specified dead time interval for the PCEP session
Capabilities List	The capabilities list for the PCEP session
Address	The local IP address of the PCEP speaker
Report Path Constraints	Indicates 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	The IP address of the PCC peer
Admin State/Oper State	The administrative and operational states of the PCC peer
Oper Keepalive/Oper Dead Timer	The operational keepalive and dead timer intervals of the PCC peer

5 ANYsec

Nokia ANYsec uses the IEEE 802.1AE (MACsec) standard in its encryption engine to encrypt MPLS payloads, while leaving the MPLS labels in clear text and unauthenticated. After the MPLS payload is encrypted, the MPLS packet is switched through an MPLS network and eventually is decrypted using the decryption engine at the terminating PE. Having the MPLS labels in clear text and unauthenticated allows any LSR router to switch the ANYsec packets from the iLER ANYsec PE to the eLER ANYsec PE. Any LSR router, including third-party routers, can manipulate the MPLS header. This includes performing label swap, pop, and push. The following figure shows the ANYsec MPLS encryption and decryption process.

Figure 23: ANYsec MPLS encryption and decryption



sw1421

5.1 ANYsec overview

Nokia ANYsec uses the IEEE 802.1AE (MACsec) standard in its datapath encryption engine and IEEE 802.1x (dot1x/MKA) as its control plane signaling protocol. ANYsec uses the user-configured MACsec connectivity association (CA) with a set of pre-shared keys (PSKs), the same way MACsec does, to negotiate datapath keys between two or more peers.

The following is a high-level overview of the ANYsec encryption process.


1. The user configures the ANYsec-specific CA, which supports the exclusive use of the CA with ANYsec encryption.

```
config>macsec>connectivity-association>anysec
```

- 2. The MACsec key agreement (MKA) key server generates the datapath security association keys (SAKs) locally on the 7705 SAR. The user configures the key-server selection priority using the **mka-key-server-priority** command.
- 3. The MKA uses CMAC-AES-128/256 to encrypt the SAKs using PSK and distributes the SAKs.
- 4. ANYsec modifies the MKA for transport over IP/UDP (described in more detail in later topics). MKA distributes the SAKs between peers, and ANYsec uses the SAKs to encrypt and decrypt the MPLS payload via the IEEE 802.1AE standard using the encryption engine.

The following figure shows a high-level implementation of ANYsec in a network.


Figure 24: ANYsec high-level implementation



5.2 ANYsec packet format

ANYsec MPLS provides packet encryption formats that are different from the MACsec packet formats. MACsec leaves the MAC 802.1AE header in clear text while authenticating the entire packet in the LAN and MACsec WAN mode leaves the dot1q VLANs in clear text and unauthenticated. With ANYsec, the entire MPLS label stack is in clear text and unauthenticated. This allows an LSR to manipulate the MPLS label stack while the system encrypts and authenticates the payload. The following figure shows ANYsec MPLS packet authentication.

Figure 25: ANYsec MPLS packet authentication



5.3 ANYsec encryption

ANYsec uses the MACsec encryption algorithm to encrypt multiple layers of OSI.

ANYsec encryption supports MACsec encryption algorithms, MPLS protocols, and per-flow encryption.

5.3.1 Encryption algorithms

ANYsec supports the following MACsec encryption algorithms:

- CMAC-AES-128/AES-256 for pre-shared key encryption used by the MKA protocol
- GCM-AES 128/256 for datapath encryption, using extended packet number (XPN) only

5.3.2 MPLS protocol support

ANYsec encrypts LSPs that are signaled by the following segment routing (SR) protocols:

- SR-ISIS for IPv4 and IPv6 (MT 0 only)
- SR-OSPF

ANYsec encrypts only the service label. The entropy label (EL) and entropy label indicator (ELI) are inserted in SR-ISIS and SR-OSPF tunnels in clear text and unencrypted when the **allow-entropy-label** command is enabled. Otherwise the EL and ELI are not inserted at all. See [ANYsec and LAG and ECMP interaction](#) for more information on ANYsec support for entropy labels.

5.3.3 Per-flow encryption

ANYsec uses a different encryption key for each encrypted LSP, thereby ensuring that all LSPs are encrypted with a unique SAK. The ANYsec MKA key server distributes the SAK to the peer. The user configures the key server selection identically to the MACsec key server selection, using the **config>macsec>conn-assoc>static-cak>mka-key-server-priority** command.

5.4 ANYsec and MACsec interaction

If ANYsec is enabled on network ports, MACsec can only be enabled on access ports. If MACsec is enabled on a network port, ANYsec cannot be enabled.

5.5 ANYsec and LAG and ECMP interaction


When LSPs are enabled with ANYsec on iLER, the LSPs do not hash their internal flows over multiple ECMP interfaces or over LAG members. Instead, each LSP on iLER maps to a dedicated ECMP interface or LAG member. If an ANYsec LSP is transporting multiple services or IP flows, all the services and IP flows map to a single interface in an ECMP group, or to a single member in a LAG bundle on the iLER.

However, if there are multiple ANYsec LSPs on the iLER, each ANYsec LSP and all the flows within the LSP can also hash to a different interface in an ECMP group or a different LAG member in a LAG bundle. LSRs can hash the ANYsec LSPs over an ECMP group interface or a LAG member freely with no restrictions.

ANYsec encrypting routers (encrypting iLER) support inserting the EL and ELI in the label stack in clear text and unencrypted with the **allow-entropy-label** command. When this command is disabled, ANYsec does not insert the EL or ELI, even if the IGP is configured to set them. On the iLER, the EL and ELI are not used for hashing. Instead, the tunnel or service is hashed to a single outgoing LAG member or ECMP interface.

The following figure shows the EL and ELI position in the packet for tunnel encryption.

Figure 26: ANYsec packet authentication with EL and ELI



5.6 Inter-AS and inter-area solutions

ANYsec LSPs support a single-area solution only. ANYsec does not support inter-AS or inter-area option B or option C (for example, seamless MPLS).

ANYsec supports inter-AS or inter-area option A with the following conditions:

- The termination of the ANYsec LSP is at the ABR or ASBR router.
- Another ANYsec LSP originates toward the other inter-AS or inter-area option A.

In addition, ANYsec does not support tunnels within tunnels. For example, an ANYsec LSP that is encrypted cannot be encapsulated in a second MPLS tunnel such as BGP LU, RSVP-TE, or LDP.

5.7 ANYsec implementation design


ANYsec supports the following implementation design:

- At the MKA signaling layer, ANYsec supports assigning the PSKs and the CAs to a group of segment routing (SR) tunnels. This assignment is known as encryption groups (EG) in ANYsec. Encryption groups support easy management of PSKs for ANYsec tunnels, because multiple tunnels can use the same PSK.
- MKA is established per pair of peers, thereby allowing encryption of each tunnel with its own dedicated SAK. The use of a dedicated SAK per tunnel ensures a high level of security. For ease of management, ANYsec supports using PSKs per group of tunnels. However, for maximum security, the SAK is only supported per pair of peers.

- ANYsec supports tunnel slicing using multi-instance IGP. Each slice is uniquely identified and a SAK is assigned per slice (per set of peers). ANYsec can also uniquely identify each LSP on each slice, and assign an EG to the LSPs through user configuration.
- ANYsec uses the MKA Layer 2 protocol for signaling, which uses IEEE 802.1x for encapsulation. To transport MKA over IP in a Layer 3 network, ANYsec encapsulates MKA over IP/UDP. The UDP port uniquely identifies the MKA packets to the ANYsec peers. The IP header transports the MKA packet from one ANYsec PE to another ANYsec PE.

The following figure shows the ANYsec design implementation and related requirements.

Figure 27: ANYsec encryption implementation design and requirements



5.8 ANYsec configuration guidelines

These topics describe the basic configuration requirements to implement ANYsec MPLS encryption.

5.8.1 Configuring ANYsec connectivity association and PSK

Because ANYsec uses MKA for its control plane and signaling SAKs, it also uses MACsec CAs and PSKs. The `config>macsec>conn-assoc>anysec` command creates a CA exclusively for the use of ANYsec encryption. When configured for ANYsec, a CA cannot also be used for MACsec.

ANYsec CAs support the `static-ca` and the `cipher-suite` commands. The following commands are blocked on ANYsec CAs:

- `clear-tag mode`
- `encryption-offset`
- `macsec-encrypt`
- `replay-protection`

- **replay-window-size**

Additionally, ANYsec does not support configuration of a MACsec MAC policy.

The following example displays an ANYsec CA configuration.

Example

```
A:node-2>config>macsec# info
-----
      connectivity-association "CA-1" create.
      anysec
      cipher-suite gcm-aes-xpn-256
      static-cak
      pre-shared-key 1 encryption-type aes-256-cmac create
      cak "2yzrsjg5sp7MYAnWpod+Nkn4SwXf70PMEfAMRpNh9Gu/badNTW0oYEG9Qi1ND0BW"
hash2
      ckn "11"
      exit
      mka-hello-interval 5
      exit
      no shutdown
      exit
```

5.8.2 Identifying and configuring ANYsec LSP

After configuring the CA as described in Configuring ANYsec connectivity association and PSK, the next step is to identify the LSP to encrypt. The benefit of ANYsec is the ability to easily enable encryption on an LSP at any time, without changing the LSP's service lease agreement (SLA). ANYsec achieves this by doing the following:

- ANYsec identifies an LSP to encrypt using a disjoint configuration, and encrypts it independently, allowing creation of MPLS networks with appropriate design and SLA that can be expanded to include encryption on LSPs in the future.
- ANYsec uses bidirectional configuration, even for LSPs that are unidirectional, and therefore the ANYsec configuration has to identify an outgoing LSP to encrypt and an incoming LSP to decrypt, under the same encryption group.

The following example displays a configuration to enable encryption on the egress LSP with the node SID 10.20.1.3, and decryption for the local node SID bound to 10.20.1.6. Both LSPs are on IGP instance ID 11.

Example:

```
*A:7705:Dut-F>config>anysec# info
-----
      reserved-label-block "Anysec"
      mka-over-ip
      mka-udp-port 10000
      exit
      security-termination-policies
      policy "ospfPolicy" create
      local-address 10.20.1.6
      protocol sr-ospf
      rx-must-be-encrypted
      no shutdown
      exit
      exit
      tunnel-encryption
      encryption-group "ospfEncrGroup" create
```



```
        ca-name "ospfCA"
        encryption-label 52006
        security-termination-policy "ospfPolicy"
        peer-tunnel-attributes
            protocol sr-ospf
        exit
        peer 10.20.1.3 create
            no shutdown
        exit
        no shutdown
    exit
exit
-----
*A:7705:Dut-F>config>anysec#
```


The security-termination policy identifies the decrypting LSP and the node SID on the local router. The decrypting LSPs node SID can arrive from multiple peers, and the SAKs used to decrypt these incoming LSPs are different for each peer.

The encryption group is a group of peers that use the same CA and PSK to secure the SAK over MKA. Each peer node SID and its local node SID use a different SAK to encrypt the datapath, but all these SAKs under the same encryption group are secured and signaled via the same CA and PSK.

A security termination policy can be part of multiple encryption groups, however, a peer can only be part of a single encryption group. This allows a set of bidirectional LSPs in an encryption group to use the same CA and PSK to secure and signal SAKs between them.

The following figure shows the implementation of ANYsec tunnel encryption using two encryption groups with the same security termination policy, and different CAs and peers.

Figure 28: ANYsec tunnel encryption implementation




5.8.3 Configuring ANYsec MKA

ANYsec uses MKA to distribute SAKs. MKA is part of the IEEE802.1x standard and is a Layer 2 protocol without an IP header. As an MPLS Layer 2.5 encryption protocol, ANYsec reuses MKA by encapsulating MKA in IP/UDP to distribute the SAK from one PE to another.

The user-configurable UDP port identifies the MKA packets on the router. Use the **config>anysec>mka-over-ip>mka-udp-port** command to configure the MKA UDP port. Nokia recommends reserving the UDP port for MKA for the entire network and ensuring it is not in use by any other application.

The following figure shows the ANYsec implementation using the MKA UDP port configuration.

Figure 29: ANYsec using MKA UDP port configuration




sw1430

5.8.4 Configuring ANYsec encryption SIDs

With ANYsec, each affected SR LSP is encrypted with a unique SAK. In segment routing, all peers wanting to send a packet to a local router, send the packet to the node SID of the local router. As such, the label stack that is pushed on the packet by different routers to switch the packet to the local router may be identical. This means if an encrypting node is also an LSR node for an upstream router, it would encrypt its own local packets to the downstream ANYsec router and would double encrypt the upstream ANYsec router packets that are going through it to the same downstream ANYsec router. For example, in the following figure PE2 would encrypt its own packet destined for PE3 and double encrypt the PE1 packets with same MPLS label stack destined for PE3. This double encryption means that PE1 packets are not decrypted correctly at PE3.

Figure 30: ANYsec double-encryption scenario



To solve this problem, Nokia introduced the encryption SID, which uniquely identifies the encrypting router within the network. The encryption SID is pushed by the encrypting router at the bottom of the label stack with S bit set if there is no entropy label present. If the entropy label is present, it will be at the bottom of the label stack and the S bit will be set on the entropy label instead of the encryption SID.


In the following figure, the encryption SID 2000 is assigned to PE1 encrypting node and encryption SID 2001 is assigned to PE2 encrypting node. Both PE1 and PE2 are sending packets to PE3.

The encryption SID pushed at the bottom of the label stack ensures that each encrypting node has a unique label stack and double encryption does not occur on PE2. The PE2 ANYsec Encryption Engine and datapath is programmed to only encrypt packets sourced at PE2. It will never encrypt packet coming from upstream routes.

Ensure the following conditions are met for the encryption SID used in the ANYsec configuration:

- The encryption SID for the ANYsec configuration must be uniquely assigned to each encrypting router.
- The encryption SID for the ANYsec configuration must be assigned from a reserved block of labels on the 7705 SAR.

Figure 31: ANYsec encryption SID



The following example displays a configuration of an encryption SID on an encrypting node.

```
A:node-2>config>router>mpls-labels# info
-----
sr-labels start 60000 end 120000
reserved-label-block "rlb-1" //reserved block for encryption sid
start-label 50000 end-label 59999
exit
-----
A:node-2>config>anysec# info
-----
reserved-label-block "rlb-1" //reserve blocked assigned to ANYsec
mka-over-ip
mka-udp-port 12345
exit
security-termination-policies
policy "STP-1" create
local-address 1.1.216.100
igp-instance-id 11
no shutdown
exit
exit
tunnel-encryption
encryption-group "EG-1" create
ca-name "CA-1"
encryption-label 50012
security-termination-policy "STP-1"
peer-tunnel-attributes
igp-instance-id 11
exit
peer 1.1.216.136 create
no shutdown
```

```

        exit
      no shutdown
    exit
  exit

```

5.9 ANYsec OAM MPLS support

ANYsec supports OAM with following SR tunnels:

- SR-ISIS for IPv4 and IPv6 (MT 0 only)
- SR-OSPF

ANYsec supports the use of the following in-band OAM commands:

- LSP ping
- SDP ping
- VCCV ping
- service ping
- L2 service ping




Note: In-band OAM packets are transmitted encrypted over an MPLS tunnel and out-of-band OAM packets are transmitted unencrypted. For example, ANYsec does not encrypt an ICMP reply that comes out-of-band, such as in IP/UDP.

5.10 ANYsec service encryption

The 7705 SAR can encrypt a service using ANYsec and the IEEE 802.1AE encryption engine.

The following figure shows the difference between tunnel and service encryption. With tunnel encryption, all services traversing the tunnel are encrypted. In contrast, service encryption allows the user to choose a specific service for encryption using its own specific PSK and SAK. The LSP carrying the services can serve multiple services that are in clear text or encrypted.

Figure 32: Tunnel and service encryption



sw4625

The 7705 SAR supports enabling both tunnel encryption and service encryption simultaneously on the same encrypted tunnel. This allows a service within the encrypted tunnel to be encrypted with its own specific SAK, while all other services traversing the tunnel are bulk-encrypted using a different SAK. For example, a provider can initially encrypt an entire tunnel with a single SAK, such that all traffic and services over that tunnel are encrypted using this SAK. Later, the provider might decide that one or more services traversing this encrypted tunnel need to be secured with their own specific SAKs and separated from the tunnel bulk encryption. The provider can then assign a unique encryption group to these services, forcing them to distribute their own SAKs via their dedicated MKA sessions. These services are then separated from the tunnel's bulk encryption, using their own MKA sessions for key distribution and unique SAKs for data encryption.

With service encryption, a tunnel can transport either encrypted or clear text services. When an encryption group is assigned to a service, that service becomes a secure service, and encryption is enabled for it. When encryption is enabled, the service only transmits and accepts encrypted traffic. Any clear text traffic is rejected and dropped.

ANYsec encryption is supported on the following services when the underlying protocol is SR IS-IS or SR-OSPF:

- spoke SDP
VPRN (Layer 3 spoke-SDP interface), IES (Layer 3 spoke-SDP interface), Cpipe, and Epipe (including PW-RED and PW-switching)
- autobind
VPRN BGP-IPVPN

ANYsec service encryption is only supported within a single area; it is not supported with inter-AS or interarea options B or C. ANYsec service encryption is possible with inter-AS option A, provided that encryption is terminated on the ASBR or ABR, and that encryption on the neighboring ASBR or ABR is restarted by assigning ANYsec to those services

ANYsec service encryption only supports global prefix SID ranges for segment routing.

5.10.1 Encryption Groups

Encryption Groups (EGs) can be assigned to a service to enable encryption for that service. An EG contains the CA, which contains the PSK used for the MKA to secure the SAK. As with tunnel encryption, the MKA is established for each pair of PEs in a service. The MKA distributes the SAK to each pair of PEs. Each PE pair within a service (such as VPRN) is encrypted with a different SAK, providing additional security. If a pair of PE SAKs is compromised, other PE pairs in the same service remain uncompromised.

An EG can be assigned to a single service or to multiple services. If assigned to multiple services, the services use the same MKA between peers and the same datapath key (SAK) to encrypt their packets between the peers. This provides scaling benefits for ANYsec, as multiple services are encrypted using the same SAK. The peer IP address must be configured under the EG. The IP address should match the BGP neighbor IP address or TLDP far-end IP address.

EGs function as follows:

- The provider can choose to encrypt a service or multiple services using a single MACsec CA and PSK (and its SAK), or assign each service a unique CA, PSK, and SAK. As with tunnel encryption, this creates an EG concept.
- An EG can be used for a set of services to be encrypted using a single CA and its corresponding PSK (CKN, CAK).

- Services using the same EG use the same local encryption SID, which uniquely identifies the source of encrypting.
- Each EG must be configured with a set of peers (PEs) that the service is encrypting traffic toward and receiving encrypted traffic from. For service encryption, these peers must match the peers discovered by dynamic protocols, such as BGP and TLDP, for each service. If a service discovers a peer through overlay signaling (BGP or TLDP), and this peer is not configured under the EG for that service, traffic is dropped.
- Each EG must be configured with its encryption SID (ES) to uniquely identify it as an encrypting source of the service. The peer's ES is discovered automatically through the MKA SCI field.

Depending on security needs, services can be configured with a single EG or with multiple EGs. When the appropriate EGs are configured with their corresponding peers, the EG is assigned to a service under autobind or within the context of its spoke SDP. The following applies:

- When an EG is assigned to an autobind, all LSPs used by that autobind inherit the EG configuration, including the CA, ES, and local address.
- For PW switching, each spoke SDP has its own EG. As a result, the encryption terminates on the first spoke and switches to the second spoke, where it is encrypted again.



Note: Two services that use two BGP groups with different local addresses to the same peer cannot share the same EG.


5.10.1.1 Shared EG between services

To share a single EG between multiple services, ensure the following rules are not violated:

1. The services must have the same peers.
2. If the service is VPRN it must use the same protocol (SR-ISIS or SR-OSPF) and the same algorithm.
3. The transport tunnel used for the services must be the same (if 1 and 2 are true, the same transport tunnel is used).
4. If the service is VPRN, autobind and Layer 3 spoke-SDP interfaces can have the same EG or different EGs.

The following figure shows an example of two encrypted services using the same tunnel and EG.

Figure 33: Two encrypted services both using the same tunnel and EG



sw4626


When the EG is shared between multiple services, only a single MKA session is established between the peers. Both services use that MKA session with its PSK to secure the SAK. The MKA session distributes the SAK between the peers, and both services are encrypted with the same SAK.

5.10.1.2 Unique EG between services

In cases where each service belongs to a unique customer, and each customer has its own security and key management requirements, it is possible to assign an EG per service. Each EG creates a dedicated MKA session with its own PSK securing its unique SAK.

The following figure shows an example of two encrypted services using the same tunnel but different EGs.

Figure 34: Two encrypted services using the same tunnel and different EGs



sw4627

5.10.2 Security termination policy

The security termination policy, as with tunnel encryption, contains the tunnel protocol information and IGP instance. It also includes the local IP address, which identifies the system IP address or loopback address used to identify the local tunnel. This local IP address must match the BGP source IP address or TLDP source IP address; it is the IP address that identifies the SR tunnel terminating on the router for the service using that security termination policy. Multiple EGs can have the same security termination policy.

The security termination policy is a peer and transport tunnel concept. A peer can contain multiple services, some encrypted and others clear text. When a service is encrypted, it accepts only encrypted traffic; any clear text traffic is dropped. In scenarios where the node is connected to PEs that have a mix of ANYsec tunnel encryption, service encryption, and clear text traffic, the **rx-must-be-encrypted** command must be disabled in order to attach the security termination policy to a service EG and tunnel EG.

5.10.3 Service encryption considerations

When identical EGs are assigned to multiple services resolving through the same peer (far-end IP address), the resolution must be through the same transport tunnel and resolution filter tunnel type. Additionally, when the same EG is assigned to multiple services, the same SAK is used to encrypt packets for all services that share the same EG.


Similar to tunnel encryption, ANYsec service encryption can coexist with MACsec. If ANYsec service encryption is enabled, MACsec can be configured on access ports.

5.10.4 Inter-AS and intra-AS solutions

Layer 2 and Layer 3 services can be configured in a single area or across multiple areas or ASs. Most deployments use Option A and Option C, though some also use Option B. ANYsec service encryption supports only Option A.

The following figure shows an example of an encrypted service with Option A.

Figure 35: Services Option A



In the scenario shown above, the service is decrypted at the ASBR or ABR PE, after which routes are leaked to the corresponding ASBR or ABR PE in the other domain. Traffic is forwarded in clear text, and in the second domain the PE can re-encrypt the service. The user data is in clear text between the ASBR and ABR, and between the edge PEs.


5.10.5 Key management

ANYsec uses MKA to distribute datapath keys (SAKs). For ANYsec, the MKA is established for each EG and each peer in the EG.

The MKA is uniquely identified through the source IP address, destination IP address, and SCI. The SCI is made up of the ES and endpoint ID.

The following figure shows an example of the MKA for an encrypted service.


Figure 36: MKA for encrypted service



In [Figure 36: MKA for encrypted service](#), if the VPRN on PE1 has two peers (PE3 and PE2) for EG1, two MKA sessions are created: one from PE1 to PE2 and another from PE1 to PE3, both for the same VPRN. The source IP address of the MKA is the local address in the security termination policy. The destination IP address is the peer IP address in EG1, which has two peers (PE2 and PE3). Each pair of local addresses, peer addresses, and SCIs create an MKA session.

As a result, each EG and PE within that EG is assigned a unique SAK for encryption, as shown in the following figure.

Figure 37: SAK for encrypted service



If two VPRNs between two PEs use the same EG and security termination policy, only a single MKA session is created. This is because the source IP address of the MKA is the security termination policy local address, which is common to both VPRNs; the destination IP address is the peer IP address in the EG, which is also common. In addition, the ES is common for both VPRNs. In this case, because one MKA session serves both VPRNs, the same SAK is signaled for both. In this example, each VPRN must be configured with a unique EG to have its own MKA and SAK.

5.10.6 Packet format for services

This section provides examples of packet formats for encrypted services in different scenarios.

Example 1: two services using the same EG

In this example, two PEs that have two VPRNs encrypted with the same EG over SR-MPLS.

This scenario occurs when a customer owns both services and wants to manage the CA and PSKs jointly for the two services. The following applies:

- The same EG can be used across multiple services if those services use the same BGP instance.
- If the same EG is configured for two services on the same two PEs, the following applies:
 - A single MKA session runs from the local address of the security termination policy to the EG peer IP address. That is, the MKA source IP address is the security termination policy local address assigned to that EG, and the destination IP address is the peer IP address of the same EG used by the two services.
 - Because there is a single MKA session between the two services, a single datapath encryption key (SAK) is used to encrypt traffic from both services. If the two services need to use different SAKs, each service must be assigned its own EG.
 - The ES is common between the two services because they are encrypted using the same SAK.

- From packets traveling from PE1 to PE3, both services appear identical, as indicated by the format shown in the following figures.

Figure 38: VPRN1 RT 1:1 service






Figure 39: VPRN2 RT 2:2 service



Example 2: Services and SR tunnel with different EGs

The following figure shows an example of two PEs that have two VPRNs encrypted with two different EGs over the same SR MPLS tunnel.

Figure 40: Example 2: Encrypted service and SR tunnel with different EGs



This scenario occurs when the two services are owned by different customers, and they require additional security by having separate CAs, PSKs, and SAKs.

When two separate EGs are configured for two services on the same two PEs, the following applies:

- There are two MKAs from the local address of the security termination policy of each EG to each EG peer IP address. The MKA source IP address is the local address of the security termination policy assigned to that EG, and the destination IP address is the peer IP address on the assigned EG for the two services. If the same BGP instance is used for both services, even if the EGs are different, the local address of the security termination policy and the peer IP address of the EG can still be the same. For example, the same security termination policy is used for both EGs and the same peer IP address is configured on both EGs. In this case, the unique identifier for the MKA is the SCI because each EG requires a unique ES. As a result, the unique MKA1 identifier is <srcIP= 2.1.1.1, dstIP=4.1.1.1, ES-1>, and MKA2 is <srcIP= 2.1.1.1, dstIP=4.1.1.1, ES-2>.
- Because there is a unique MKA for each service, a unique SAK is used to encrypt service user traffic for each service.
- From online packets going from PE1 to PE3, there are two unique packet signatures, as shown in the following figures.

Figure 41: VPRN1 RT 1:1 service





Figure 42: VPRN2 RT 1:1 service



5.11 ANYsec command reference

5.11.1 Command hierarchies

- [Configuration commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)
- Tools Commands (see the Tools chapter in the 7705 SAR OAM and Diagnostics Guide)

5.11.1.1 Configuration commands

```

config
- anysec
- mka-over-ip
- mka-udp-port port
- no mka-udp-port
- reserved-label-block name
- no reserved-label-block
- security-termination-policies
- policy policy-name [create]
- no policy policy-name
- igp-instance-id id
- local-address ip-address | ipv6-address

```

```

- no local-address
- protocol protocol
- [no] rx-must-be-encrypted
- [no] shutdown
- [no] service-encryption
- encryption-group group-name [create]
- no encryption-group group-name
- ca-name ca-name
- no ca-name
- encryption-label label
- no encryption-label
- peer ip-address|ipv6-address [create]
- no peer ip-address|ipv6-address
- [no] shutdown
- security-termination-policy policy
- no security-termination-policy
- [no] shutdown
- tunnel-encryption
- encryption-group group-name [create]
- no encryption-group group-name
- ca-name ca-name
- no ca-name
- [no] allow-entropy-label
- encryption-label label
- no encryption-label
- peer ip-address | ipv6-address [create]
- no peer ip-address | ipv6-address
- [no] shutdown
- peer-tunnel-attributes
- igp-instance-id id
- protocol protocol
- security-termination-policy policy
- no security-termination-policy
- [no] shutdown

```

5.11.1.2 Show commands

```

show
- [no] anysec
- mka-over-ip
- tunnel-encryption encryption-group group-name [peer ip-address | ipv6-address]
[detail] [mka-session]
- tunnel-encryption encryption-group group-name [peer ip-address | ipv6-address]
[statistics] [mka-session]
- tunnel-encryption [detail] [mka-session]
- tunnel-encryption [statistics] [mka-session]

```

5.11.1.3 Clear commands

```

clear
- anysec encryption-group group-name peer ip-address | ipv6-address statistics

```

5.11.1.4 Debug commands

```

debug

```



```
- anysec
  - encryption-group group-name peer ip-address | ipv6-address [detail]
  - peer ip-address | ipv6-address [detail]
```

5.11.2 Command descriptions

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

5.11.2.1 Configuration commands

anysec

Syntax

anysec

Context

config

Description

This command enables the context to configure ANYsec commands.

Default

n/a

mka-over-ip

Syntax

mka-over-ip

Context

config>anysec

Description

This command enables the context to configure the MACsec Key Agreement (MKA) over IP.

Default

n/a

mka-udp-port

Syntax

mka-udp-port *port*

no mka-udp-port

Context

config>anysec>mka-over-ip

Description

This command configures the UDP port that identifies the MKA packet on the system.

Nokia recommends configuring this UDP port network wide. In addition, ensure the UDP port is not used by any other protocols in the network.

The **no** form of this command removes the configured UDP port.

Default

no mka-udp-port

Parameters

port

the UDP port that identifies the MKA packet

Values 1024 to 49151

reserved-label-block

Syntax

reserved-label-block *name*

no reserved-label-block

Context

config>anysec

Description

This command assigns an existing label block reserved for ANYsec encryption SIDs. The label block must have been configured with the **config>router>mpls-labels>reserved-label-block** command. Without the reserved label block, ANYsec cannot assign encryption SIDs.

The encryption SID uniquely identifies the encrypting node within a network and avoids double encryption scenarios. The encryption SID can be assigned per encryption group. Each encryption group must have a unique encryption SID.

To save label space, Nokia recommends limiting the number of encryption SIDs within a network.

The **no** form of this command removes the reserved label block.

Default

no reserved-label-block

Parameters

name

specifies the name of an existing reserved label block to use

security-termination-policies

Syntax

security-termination-policies

Context

config>anysec

Description

This command enables the context to configure local security termination policies.

Default

n/a

policy

Syntax

policy *policy-name* [**create**]

no policy *policy-name*

Context

config>anysec>sec-term-pols

Description

This command configures an ANYsec security termination policy for locally terminating tunnels.

The **no** form of this command removes the security termination policy.

Default

no policy

Parameters

policy-name

the local security termination policy name, up to 32 characters

create

keyword required when first creating the policy. When the policy is created, you can navigate into the context without the **create** keyword.

igp-instance-id

Syntax

igp-instance-id *id*

Context

config>anysec>sec-term-pols>policy

config>anysec>tnl-enc>enc-grp>peer-tnl-attrs

Description

This command configures the IGP instance ID.

For security termination policies, the IGP instance ID must match the IGP instance that the incoming encrypted LSP was signaled on.

For tunnel encryption, IGP instance ID must be a match for the outgoing tunnel. This IGP instance ID should match the IGP instance the LSP is being signaled on in order for ANYsec to encrypt the LSP.

Default

igp-instance-id 0

Parameters

id

the IGP instance used to tunnel packets for the session.

Values 0 to 31

local-address

Syntax

local-address *ip-address* | *ipv6-address*

no local-address

Context

config>anysec>sec-term-pols>policy

Description

This command configures the local IPv4 or IPv6 address for the system IP or loopback node SID. This is used to identify the LSP as an incoming ANYsec tunnel and will configure the datapath for decryption of this LSP.

The **no** form of this command removes the IP address that is associated with ANYsec decryption.

Default

no local-address

Parameters

ip-address

the IPv4 address of the local address

ipv6-address

the IPv6 address of the local address

protocol

Syntax

protocol *protocol*

Context

config>anysec>sec-term-pols>policy

Description

This command configures a routing protocol that is used to advertise the node SID of the incoming tunnel.

Default

protocol sr-isis

Parameters

protocol

the protocol name to match on

Values sr-isis, sr-ospf

rx-must-be-encrypted

Syntax

[no] **rx-must-be-encrypted**

Context

config>anysec>sec-term-pols>policy

Description

This command specifies whether all traffic that is not secured using ANYsec that is received on the port is dropped.

When the command is enabled, all arriving traffic that is not secured using MACsec is dropped.

The **no** form of the command accepts all traffic received on the port whether or not it is secured by ANYsec.

Default

no rx-must-be-encrypted

shutdown**Syntax**

[no] shutdown

Context

config>anysec>sec-term-pols>policy

Description

This command places the tunnel in an operationally down state, causing traffic to drop but leaving ANYsec enabled.

The **no** form of this command places the tunnel in an operationally up state.

Default

shutdown

service-encryption**Syntax**

service-encryption

Context

config>anysec

Description

This command enables the context to configure ANYsec for service encryption.

Default

n/a

tunnel-encryption

Syntax

tunnel-encryption

Context

config>anysec

Description

This command enables the context to configure ANYsec for tunnel encryption.

Default

n/a

encryption-group

Syntax

encryption-group *group-name* [create]

no encryption-group *group-name*

Context

config>anysec>srv-enc

config>anysec>tnl-enc

Description

This command creates an encryption group for service or tunnel encryption.

An encryption group for service encryption is a group of tunnels that use the same CA and its PSK.

An encryption group for tunnel encryption is a group of LSPs that use the same CA and the same PSKs. The PSK is used to secure the SAK for distribution to other peers.



Note: Although the LSPs are unidirectional, ANYsec is a bidirectional concept where a pair of LSPs between two peers are encrypted and decrypted. Each pair of LSP uses its own SAK for maximum security, though they may share the same CA and PSK with all other LSPs in the encryption group to secure the SAK.

The **no** form of the command removes the encryption group.

Default

no encryption-group

Parameters

group-name

the encryption group name, up to 32 characters

create

keyword required when first creating the encryption group. When the encryption group is created, you can navigate into the context without the **create** keyword.

allow-entropy-label

Syntax

[no] **allow-entropy-label**

Context

config>anysec>tnl-enc>enc-grp

Description

This command enables the ANYsec encryption group to support entropy labels. When enabled, the EL and ELI are inserted in the label stack in clear text.

The **no** form of this command disables entropy labels for the ANYsec encryption group. The EL and ELI are not inserted in the label stack, even when the IGP is configured to set them.

Default

no allow-entropy-label

ca-name

Syntax

ca-name *ca-name*

no ca-name

Context

config>anysec>svc-enc>enc-grp

config>anysec>tnl-enc>enc-grp

Description

This command assigns an existing CA to be used for this encryption group. The CA must have been configured for ANYsec with the **config>macsec>connectivity-association>anysec** command.

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

Default

no ca-name

Parameters

ca-name

specifies the name of an existing ANYsec CA to use

encryption-label

Syntax

encryption-label *label*

no encryption-label

Context

config>anysec>svc-enc>enc-grp

config>anysec>tnl-enc>enc-grp

Description

This command creates an encryption group label ID that identifies the encrypting node or encryption group used to encrypt a tunnel.

The encryption SID uniquely identifies the encrypting node within a network to avoid double encryption scenarios. The encryption SID can be assigned per encryption group. Each encryption group must have a unique encryption SID. To save label space, Nokia recommends limiting the number of encryption SIDs within a network. To configure the encryption SID, a **reserved-label-block** command must be configured under the **anysec** context. The encryption SID is programmed at the bottom of the stack with S bit set if there is no entropy label present. If the entropy label is present, it will be at the bottom of the label stack and the S bit will be set on the entropy label instead of the encryption SID.

The **no** form of this command removes the encryption group label.

Default

no encryption-label

Parameters

label

specifies the encryption group label ID

Values 0 to 1048575

peer

Syntax

peer *ip-address* | *ipv6-address* [**create**]

no peer *ip-address* | *ipv6-address*

Context

```
config>anysec>svc-enc>enc-grp
```

Description

This command configures peer address as part of the specified encryption group..

Peers are identified by their unique tunnel attributes. A peer can only be part of a single encryption group and cannot participate in multiple encryption groups.

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

Default

no peer

Parameters

ip-address

the IPv4 address of the peer

ipv6-address

the IPv6 address of the peer

create

keyword required when first configuring the peer's address. When the peer is configured, you can navigate into the context without the **create** keyword.

peer

Syntax

peer *ip-address* | *ipv6-address* [**create**]

no peer *ip-address* | *ipv6-address*

Context

```
config>anysec>tnl-enc>enc-grp
```

Description

This command configures the IPv4 or IPv6 address of the peer's node SID that is part of this encryption group.

This configuration identifies the peer's segment routing node SID and programs the datapath for encryption of the LSP. When the label stack is downloaded, the encryption SID is also included at the bottom of the stack with the S bit set if there is no entropy label present. If the entropy label is present, it will be at the bottom of the label stack and the S bit will be set on the entropy label instead of the encryption SID.

The **no** form of this command removes the peer's node SID from the ANYsec configuration. Therefore, the LSP is not encrypted and all the traffic is transmitted in clear text.

Default

no peer

Parameters

ip-address

the IPv4 address of the peer's node SID

ipv6-address

the IPv6 address of the peer's node SID

create

keyword required when first configuring the IPv4 or IPv6 address of the peer's node SID. When the peer is configured, you can navigate into the context without the **create** keyword.

shutdown

Syntax

[no] shutdown

Context

config>anysec>svc-enc>enc-grp>peer

config>anysec>tnl-enc>enc-grp>peer

Description

This command shuts down the ANYsec and encryption to the peer.

When ANYsec is shutdown to the peer, clear traffic is forwarded on the LSP to the peer.

The **no** form of this command enables ANYsec and encryption to the peer.

Default

shutdown

peer-tunnel-attributes

Syntax

peer-tunnel-attributes

Context

config>anysec>tnl-enc>enc-grp

Description

This command enters the context to configure the peer tunnel attributes.

Tunnel attributes are used to match and identify the outgoing tunnels for encryption with ANYsec. A single tunnel attribute is used for multiple peers. Because an LSP is unidirectional, the outgoing tunnel can have different attributes from the incoming tunnel (for example, security termination policy).

Default

n/a

protocol

Syntax**protocol** *protocol***Context**

config>anysec>tnl-enc>enc-grp>peer-tnl-attrs

Description

This command configures a routing protocol as a match criterion for the outgoing tunnel.



Note: Because the LSP is unidirectional, the incoming tunnel protocol may be different from the outgoing tunnel protocol.

Default

protocol sr-isis

Parameters*protocol*

the protocol name to match on

Values sr-isis, sr-ospf

security-termination-policy

Syntax**security-termination-policy** *policy***no security-termination-policy****Context**

config>anysec>svc-enc>enc-grp

config>anysec>tnl-enc>enc-grp

Description

This command configures the security termination policy that identifies the set of attributes for locally terminating tunnels for the specified encryption group.

The security termination policy for tunnel encryption contains the tunnel protocol information and IGP instance in addition to the local IP address, which identifies the system IP or loopback IP used to identify the local underlying tunnel for the encrypted service. This local IP must match the BGP source IP or TLDP source IP and is the IP address that identifies the SR tunnel terminating on the router for the service

that is using the specified security termination policy. Multiple encryption groups can have the same security termination policy.

The **no** form of this command removes the security termination policy.

Default

no security-termination-policy

Parameters

policy

the local security termination policy name, up to 32 characters

shutdown

Syntax

[no] shutdown

Context

config>anysec>svc-enc>enc-grp

config>anysec>tnl-enc>enc-grp

Description

This command shuts down ANYsec on the encryption group.

When ANYsec is shut down on a tunnel , clear traffic is forwarded on all the LSPs in the encryption group.

The **no** form of this command enables ANYsec and encryption for the encryption group.

Default

shutdown

5.11.2.2 Show commands

mka-over-ip

Syntax

mka-over-ip

Context

show>anysec

Description

This command displays MKA-over-IP information.

Output

The following output is an example of MKA over IP information, and [Table 41: MKA over IP field descriptions](#) describes the fields.

Output example

```
*A: Sar18 Dut-B>show>anysec# mka-over-ip#
=====
MKA over IP
=====
MKA over IP session UDP Port: 12345
Reserved Label Block : (Not Specified)
Operational Status : in-service
=====
```

Table 41: MKA over IP field descriptions

Label	Description
MKA over IP session UDP Port	The MKA-over-IP UDP port number
Reserved Label Block	The reserved label block
Operational Status	The operational status

tunnel-encryption

Syntax

```
tunnel-encryption encryption-group group-name [peer ip-address | ipv6-address] [detail] [mka-session]
tunnel-encryption encryption-group group-name [peer ip-address | ipv6-address] [statistics] [mka-session]
tunnel-encryption [detail] [mka-session]
tunnel-encryption [statistics] [mka-session]
```

Context

```
show>anysec
```

Description

This command displays ANYsec tunnel encryption information.

Parameters

- group-name*
specifies the encryption group name, up to 32 characters
- ip-address*
specifies the IPv4 address

ipv6-address

specifies the IPv6 address

detail

keyword to specify the detailed tunnel encryption group information

mka-session

keyword to specify MKA session information

statistics

keyword to specify MKA statistical information

Output

The following output is an example of tunnel encryption information, and [Table 42: Tunnel encryption field descriptions](#) describes the fields.

Output example - detail

```
*A:7705:Dut-C# show anysec tunnel-encryption encryption-group "isisEncrGroup" detail

=====
Encryption Group: isisEncrGroup
=====
Group Admin State      : Up
CA Name                : isisCA
Local Encryption Label : 50003
Security Termination Policy: isisPolicy
Policy Admin State     : Up
Policy Protocol        : sr-isis
Policy IGP Instance ID : 0
Policy FlexAlgo ID     : 0
Local Address          : 10.20.1.3
Rx Must Be Encrypted   : Yes
Terminating Label      : 22003

-----
Peer 10.20.1.6 configuration
-----
Peer Admin State      : Up
Peer Protocol         : sr-isis
Peer IGP Instance ID  : 0
Peer FlexAlgo ID      : 0
Peer Tunnel ID        : 524295
Peer Encryption Label  : 50006
Last Inconsistent Rx SCI: 0000000000000000
Peer Oper State       : Down
Peer Oper Down Reason  : mka-down
Peer Enhanced Padding  : disabled

-----
SC stats for peer 10.20.1.6
-----

-----
txSCSecyStats
-----
Tx SCI                 : 00000000c3530001
Encrypted Packets      : 1653
Encrypted Octets       : 674424
Enqueue Errors         : 0
```

```
Other Errors          : 0
-----
rxSCSecyStats
-----
Rx SCI                : 00000000c3560001
Delay Packets         : 0
Not Valid Packets     : 0
Unchecked Packets     : 0
OK Packets           : 0
Decrypted Octets      : 0
Enqueue Errors       : 0
Other Errors         : 0
=====
*A:7705:Dut-C#
```

Output example - statistics

```
*A:7705:Dut-C# show anysec tunnel-encryption encryption-group "isisEncrGroup" statistics
=====
Encryption Group: isisEncrGroup
=====

-----
SC stats for peer 10.20.1.6
-----

txSCSecyStats
-----
Tx SCI                : 00000000c3530001
Encrypted Packets     : 1653
Encrypted Octets      : 674424
Enqueue Errors       : 0
Other Errors         : 0
-----

rxSCSecyStats
-----
Rx SCI                : 00000000c3560001
Delay Packets         : 0
Not Valid Packets     : 0
Unchecked Packets     : 0
OK Packets           : 0
Decrypted Octets      : 0
Enqueue Errors       : 0
Other Errors         : 0
=====
```

Table 42: Tunnel encryption field descriptions

Label	Description
Group Admin State	The encryption group administrative state
CA Name	The connectivity association name

Label	Description
Local Encryption Label	The local encryption label
Security Termination Policy	The security termination policy name
Policy Admin State	The policy administrative state
Policy Protocol	The policy protocol
Policy IGP Instance ID	The policy IGP instance ID
Policy FlexAlgo ID	The policy flexible algorithm ID
Local Address	The local IP address
Rx Must Be Encrypted	Indicates all arriving traffic must be encrypted or is discarded
Terminating Label	The terminating label
Peer <IP address> configuration	
Peer Admin State	The peer administrative state
Peer Protocol	The peer protocol
Peer IGP Instance ID	The peer IGP instance ID
Peer FlexAlgo ID	The peer Flex Algo ID
Peer Tunnel ID	The peer tunnel ID
Peer Encryption Label	The peer encryption label
Last Inconsistent Rx SCI	The last inconsistent receive SCI value
Peer Oper State	The operational state of the peer
Peer Oper Down Reason	The reason the peer is operationally down
Peer Enhanced Padding	Indicates whether ANYsec enhanced padding is enabled or disabled
txSCSeyStats	
Tx SCI	The transmit SCI value
Encrypted Packets	The number of encrypted packets
Encrypted Octets	The number of encrypted octets
rxSCSeyStats	
Rx SCI	The receive SCI value
Delayed Packets	The number of delayed packets
Not Valid Packets	The number of invalid packets
Unchecked Packets	The number of unchecked packets

Label	Description
OK Packets	The number of OK packets
Decrypted Octets	The number of decrypted octets

5.11.2.3 Clear commands

anysec

Syntax

anysec encryption-group *group-name* **peer** *ip-address* | *ipv6-address* **statistics**

Context

clear

Description

This command clears ANYsec statistics.

Default

n/a

Parameters

group-name

the encryption group name, up to 32 characters

ip-address

the IPv4 address of the peer

ipv6-address

the IPv6 address of the peer

statistic

keyword used to clear ANYsec statistics

5.11.2.4 Debug commands

anysec

Syntax

[no] anysec

Context

debug

Description

This command enables ANYsec debugging.

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

Default

no anysec

encryption-group**Syntax**

encryption-group *group-name* **peer** *ip-address* | *ipv6-address* [**detail**]

Context

debug>anysec

Description

This command debugs ANYsec encryption groups.

Default

n/a

Parameters

group-name

the encryption group name, up to 32 characters

ip-address

specifies a peer IPv4 address

ipv6-address

specifies a peer IPv6 address

detail

keyword used to specify detailed information

peer**Syntax**

peer *ip-address* | *ipv6-address* [**detail**]

Context

debug>anysec

Description

This command debugs ANYsec peers.

Default

n/a

Parameters

ip-address

specifies a peer IPv4 address

ipv6-address

specifies a peer IPv6 address

detail

keyword used to specify detailed information

6 Label Distribution Protocol

This chapter provides information to enable the Label Distribution Protocol (LDP).

Topics in this chapter include:

- [Label Distribution Protocol](#)
- [LDP point-to-multipoint support](#)
- [MLDP fast upstream switchover](#)
- [LDP IPv6](#)
- [LDP process overview](#)
- [Configuration notes](#)
- [Configuring LDP with CLI](#)
- [LDP command reference](#)

6.1 Label Distribution Protocol

Label Distribution Protocol (LDP) is used to distribute labels in non-traffic-engineered applications. LDP allows routers to establish LSPs through a network by mapping network-layer routing information directly to data link LSPs.

An LSP is defined by the set of labels from the ingress LER to the egress LER. 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 ingress LER 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 by each LSR, where each LSR splices incoming labels for the FEC to the outgoing label assigned to the next hop for the 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 (DU). For LDP on the 7705 SAR, Downstream Unsolicited (DU) mode is implemented.

This section contains the following topics:

- [LDP and MPLS](#)
- [LDP architecture](#)
- [LDP subsystem interrelationships](#)
- [Execution flow](#)
- [Label exchange](#)

- LDP filters
- LDP FEC statistics
- Multi-area and multi-instance extensions to LDP
- ECMP support for LDP
- Graceful restart helper
- Graceful handling of resource exhaustion
- LDP support for unnumbered interfaces
- LDP FRR
- LDP-to-segment routing stitching for IPv4 /32 prefixes (IS-IS)
- LDP FRR remote LFA and TI-LFA backup using an SR tunnel for IPv4 /32 prefixes (IS-IS)
- TCP MD5 and enhanced TCP authentication

6.1.1 LDP and MPLS

LDP performs dynamic label distribution in MPLS environments. The LDP operation begins with a Hello discovery process network to form an adjacency with an LDP peer in the network. LDP peers are two MPLS routers 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 and to distribute its own label information (LDP is bidirectional), and 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 pseudowires (PWs) 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 is 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 that enable fast access and packet identification are maintained in the forwarding plane.

After an unlabeled packet ingresses the 7705 SAR, classification policies associate it with a FEC, the appropriate label is imposed on the packet, and then the packet is forwarded. Other actions can also take place on a packet before it is forwarded, including imposing additional labels, other encapsulations, or learning actions. After 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 performed on the packet and then the packet is forwarded.

The LDP implementation provides support for DU, ordered control, and liberal label retention mode.

For LDP label advertisement, DU mode is supported. To prevent filling the uplink bandwidth with unassigned label information, Ordered Label Distribution Control mode is supported.

A PW/VLL label can be dynamically assigned by targeted LDP operations. Targeted LDP allows the inner labels (that is, the VLL labels) in the MPLS headers to be managed automatically. This makes it easier for operators to manage the VLL connections. There is, however, additional signaling and processing overhead associated with this targeted LDP dynamic label assignment.

6.1.1.1 BFD for T-LDP

BFD is a simple protocol for detecting failures in a network. BFD uses a "hello" mechanism that sends control messages periodically to the far end and receives periodic control messages from the far end. BFD is implemented in asynchronous mode only, meaning that neither end responds to control messages; rather, the messages are sent in the time period configured at each end.

A T-LDP session is a session between either directly or non-directly connected peers and requires that adjacencies be created between two peers. BFD for T-LDP sessions allows support for tracking of failures of nodes that are not directly connected. BFD timers must be configured under the system router interface context before being enabled under T-LDP.

BFD tracking of an LDP session associated with a T-LDP adjacency allows for faster detection of the status of the session by registering the loopback address of the peer as the transport address.

6.1.2 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 43: LDP 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 interprocess communication through the event queue, as well as through updates to the various data structures. The following list describes the primary data structures that LDP maintains:

- FEC/label database – this database contains all the FEC-to-label mappings, including both sent and received. It also contains both address FECs (prefixes and host addresses) as well as service FECs (L2 VLLs).
- 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

6.1.3 LDP subsystem interrelationships

[Figure 43: LDP subsystem interrelationships](#) shows the relationships between LDP subsystems and other 7705 SAR subsystems. The following sections describe how the subsystems work to provide services.

6.1.3.1 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 large chunks and managing the memory internally to LDP. When LDP is shut down, it releases all memory allocated to it.

6.1.3.2 Label manager

LDP assumes that the label manager is up and running. LDP aborts initialization if the label manager is not running. The label manager is initialized at system boot-up; therefore anything that causes it to fail likely indicates that the system is not functional. The 7705 SAR uses a label range from 28 672 (28K) to 131 071 (128K-1) to allocate all dynamic labels, including VC labels.

Figure 43: LDP subsystem interrelationships



6.1.3.3 LDP configuration

The 7705 SAR uses 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, targeted LDP sessions can be set up to specific endpoints. Targeted session parameters are configurable.

6.1.3.4 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. See the *7705 SAR System Management Guide* for logger configuration information.

6.1.3.5 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 the service manager of relevant LDP events, such as connection setups and failures, timeouts, and labels signaled or withdrawn.

6.1.4 Execution flow

LDP activity in the 7705 SAR is limited to service-related signaling. Therefore, the configurable parameters are restricted to system-wide parameters, such as hello and keepalive timeouts.

6.1.4.1 Initialization

MPLS must be enabled when LDP is initialized. LDP makes sure that the various prerequisites are met, such as ensuring that the system IP interface and the label manager are operational, and ensuring that there is memory available. It then allocates a pool of memory to itself and initializes its databases.

6.1.4.2 Session lifetime

In order for a targeted LDP session to be established, an adjacency has to be created. The LDP extended discovery mechanism requires hello messages to be exchanged between two peers for session establishment. After the adjacency is established, session setup is attempted.

6.1.4.2.1 Adjacency establishment

In the 7705 SAR, adjacency management is done through the establishment of a Service Destination Point (SDP) object, which is a service entity in the Nokia service model.

The service model uses logical entities that interact to provide a service. The service model requires the service provider to create and configure four main entities:

- customers
- services
- Service Access Points (SAPs) on local 7705 SAR routers
- SDPs that connect to one or more remote 7705 SAR routers or 77x0 SR routers

An SDP is the network-side termination point for a tunnel to a remote 7705 SAR or 77x0 SR router. An SDP defines a local entity that includes the system IP address of the remote 7705 SAR routers and 77x0 SR routers, and a path type.

Each SDP comprises:

- the SDP ID
- the transport encapsulation type, MPLS
- the far-end system IP address

If the SDP is identified as using LDP signaling, then an LDP extended hello adjacency is attempted.

If another SDP is created to the same remote destination and if LDP signaling is enabled, no further action is taken, because only one adjacency and one LDP session exists between the pair of nodes.

An SDP is a unidirectional object, so a pair of SDPs pointing at each other must be configured in order for an LDP adjacency to be established. After an adjacency is established, it is maintained through periodic hello messages.

6.1.4.2.2 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 liveness.

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 reattempted as the back-pressure eases.

6.1.5 Label exchange

Label exchange is initiated by the service manager. After an SDP is attached to a service (that is, after 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.

6.1.5.1 Implicit null label

The implicit null label option enables an eLER to receive MPLS packets from the previous-hop LSR without the outer LSP label.

The implicit null label is signaled by the eLER to the previous-hop LSR during FEC signaling by the LDP control protocol. When the implicit null label is signaled to the LSR, it pops the outer label before sending the MPLS packet to the eLER; this is known as penultimate hop popping.

The implicit null label option can be enabled for all LDP FECs for which the router is the eLER by using the **implicit-null-label** command in the **config>router>ldp** context.

If the implicit null configuration is changed, LDP withdraws all the FECs and readvertises them using the new label value.

6.1.5.2 Other reasons for label actions

Label actions can also occur for the following reasons:

- MTU changes – LDP withdraws the previously assigned label and resignals the FEC with the new Maximum Transmission Unit (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, the received label is released
- VC type unsupported – when an unsupported VC type is received, the received label is released

6.1.5.3 Cleanup

LDP closes all sockets, frees all memory, and shuts down all its tasks when it is deleted, so that it uses no memory (0 bytes) when it is not running.

6.1.6 LDP filters

The 7705 SAR supports both inbound and outbound LDP label binding filtering.

Inbound filtering (import policy) allows the user to configure a policy to control the label bindings a label switching router (LSR) accepts from its peers.

Import policy label bindings can be filtered based on the following:

- neighbor – match on bindings received from the specified peer
- prefix-list – match on bindings with the specified prefix/prefixes

The default import behavior is to accept all FECs received from peers.

Outbound filtering (export policy) allows the user to configure a policy to control the set of LDP label bindings advertised by the LSR.

Because the default behavior is to originate label bindings for the system IP address only, when a non-default loopback address is used as the transport address, the 7705 SAR does not advertise the loopback FEC automatically. With LDP export policy, the user is now able to explicitly export the loopback address to advertise the loopback address label and allow the node to be reached by other network elements.

Export policy label bindings can be filtered based on the following:

- all – all local subnets by specifying "direct" as the match protocol
- prefix-list – match on bindings with the specified prefix/prefixes



Note: In order for the 7705 SAR to consider a received label to be active, there must be an exact match to the FEC advertised together with the label found in the routing table, or a longest prefix match (if the aggregate-prefix-match option is enabled; see [Multi-area and multi-instance extensions to LDP](#)). This can be achieved by configuring a static route pointing to the prefix encoded in the FEC.

6.1.6.1 Per-LDP peer export policies

Prefix export policies provide a way to control which FEC prefixes received from other LDP peers are redistributed to a specific LDP peer. By default, all FEC prefixes are exported to the LDP peer.

FEC prefix export policies are configured using the **config>router>ldp>session-params>peer>export-prefixes** *policy-name* command.

6.1.7 LDP FEC statistics

LDP FEC statistics allow operators to monitor traffic being forwarded between any two PE routers and for all services using an LDP SDP. LDP FEC statistics are available for the egress data path at the ingress LER and LSR. Because an ingress LER is also potentially an LSR for an LDP FEC, combined egress data path statistics are provided whenever applicable. For more information, see [RSVP LSP and LDP FEC statistics](#).

6.1.8 Multi-area and multi-instance extensions to LDP

When a network has two or more IGP areas, or instances, inter-area LSPs are required for MPLS connectivity between the PE devices that are located in the distinct IGP areas. To extend LDP across multiple areas of an IGP instance or across multiple IGP instances, the current standard LDP implementation based on RFC 3036, *LDP Specification*, requires that all /32 prefixes of PEs be leaked between the areas or instances. IGP route leaking is the distribution of the PE loopback addresses across area boundaries. An exact match of the prefix in the routing table (RIB) is required to install the prefix binding in the FIB and set up the LSP.

This behavior is the default behavior for the 7705 SAR when it is configured as an Area Border Router (ABR). However, exact prefix matching causes performance issues for the convergence of IGP on routers deployed in networks where the number of PE nodes scales to thousands of nodes. Exact prefix matching requires the RIB and FIB to contain the IP addresses maintained by every LSR in the domain and requires redistribution of a large number of addresses by the ABRs. Security is a potential issue as well, as host routes leaked between areas can be used in DoS and DDoS attacks and spoofing attacks.

To prevent these performance and security issues, the 7705 SAR can be configured for an optional behavior in which LDP installs a prefix binding in the LDP FIB by performing a longest prefix match with an aggregate prefix in the routing table (RIB). This behavior is described in RFC 5283, *LDP Extension for Inter-Area Label Switched Paths*. The LDP prefix binding continues to be advertised on a per-individual /32 prefix basis.

When the longest prefix match option is enabled and an LSR receives a FEC-label binding from an LDP neighbor for a prefix-address FEC element, FEC1, it installs the binding in the LDP FIB if:

- the routing table (RIB) contains an entry that matches FEC1. Matching can either be a longest IP match of the FEC prefix or an exact match.
- the advertising LDP neighbor is the next hop to reach FEC1

When the FEC-label binding has been installed in the LDP FIB, LDP programs an NHLFE entry in the egress data path 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 RIB, 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, this may mean that the LSR used as the next hop changes; if so, the NHLFE entry for that FEC must be changed.

When a prefix is removed from the RIB, 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, LDP must use it. This may mean that the LSR used as the next hop changes; if so, the NHLFE entry for that FEC must be changed. If another match does not exist, the LSR removes the FEC binding and sends a label withdraw message to its LDP neighbors.

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 entry for the FEC elements.

6.1.9 ECMP support for LDP

Equal-Cost Multipath Protocol (ECMP) support for LDP performs load balancing for services that use LDP-based LSPs as transport tunnels, by having multiple equal-cost outgoing next hops for an IP prefix.

ECMP for LDP load-balances traffic across all equal-cost links based on the output of the hashing algorithm using the allowed inputs, based on the service type. For additional information, see "LAG and ECMP Hashing" in the 7705 SAR Interface Configuration Guide.

There is only one next-hop peer for a network link. To offer protection from a network link or next-hop peer failure, multiple network links can be configured to connect to different next-hop peers, or multiple links to the same peer. For example, an MLPPP link and an Ethernet link can be connected to two peers, or two Ethernet links can be connected to the same peer. ECMP occurs when the cost of each link reaching a target IP prefix is equal.

The 7705 SAR uses a liberal label retention mode, which retains all labels for an IP prefix from all next-hop peers. A 7705 SAR acting as an LSR load-balances the MPLS traffic over multiple links using a hashing algorithm.

The 7705 SAR supports the following optional fields as hash inputs and supports profiles for various combinations:

- hashing algorithms
 - label-only option: hashing is done on the MPLS label stack, up to a maximum of 10 labels (default)
 - label-IP option: hashing is done on the MPLS label stack and the IPv4 source and destination IP address if an IPv4 header is present after the MPLS labels
 - Layer 4 header (source or destination UDP or TCP port number) and TEID: hashing is done on the MPLS label stack, the IPv4 source and destination IP address (if present), then on the Layer 4 source and destination UDP or TCP port fields (if present) and the TEID in the GTP header (if present)
- label stack profile options on significance of the bottom-of-stack label (VC label)
 - profile 1: favors better load balancing for pseudowires when the VC label distribution is contiguous (default)
 - profile 2: similar to profile 1 where the VC labels are contiguous, but provides an alternate distribution
 - profile 3: all labels have equal influence in hash key generation
- ingress LAG port at the LSR (default is disabled)

The **use-ingress-port** option, when enabled, specifies that the ingress port is used by the hashing algorithm at the LSR. This option should be enabled for ingress LAG ports because packets with the same label stack can arrive on all ports of a LAG interface. In this case, using the ingress port in the hashing algorithm results in better egress load balancing, especially for pseudowires.

The option should be disabled for LDP ECMP so that the ingress port is not used by the hashing algorithm. For ingress LDP ECMP, if the ingress port is used by the hashing algorithm, the hash distribution could be biased, especially for pseudowires.

- system IP address – hashing on the system IP address is enabled and disabled at the system level only

All of the above options can be configured with the **lsr-load-balancing** command, with the exception of the system IP address, which is configured with the **system-ip-load-balancing** command.



Note: The global IF index is no longer a hash input for LSR ECMP load balancing. It has been replaced with the **use-ingress-port** configurable option in the **lsr-load-balancing** command. As well, the default treatment of the MPLS label stack has changed to focus on the bottom-of-stack label (VC label). In previous releases, all labels had equal influence.

LSR load balancing can be configured at the system level or interface level. Configuration at the interface level overrides the system-level settings for the specific interface. Configuration must be done on the ingress network interface (that is, the interface on the LDP LSR node that the packet is received on).

Configuration of load balancing at the interface level provides some control to the user; for example, the label-IP option can be disabled on a specific interface if labeled packets received on the interface include non-IP packets that can be confused by the hash routine for IP packets. Disabling the label-IP option can be used in cases where the first nibble of a non-IP packet is a 4, which would result in the packet being hashed incorrectly if the label-IP option was enabled.

If ECMP is not enabled, the label from only one of the next-hop peers is selected and installed in the forwarding plane. In this case, the algorithm used to distribute the traffic flow looks up the route information, and selects the network link with the lowest IP address. If the selected network link or next-hop peer fails, another next-hop peer is selected, and LDP reprograms the forwarding plane to use the label sent by the newly selected peer.

ECMP is supported on all Ethernet ports in network mode, and is also supported on the 4-port OC3/STM1 Clear Channel Adapter card when it is configured for POS (ppp-auto) encapsulation and network mode.

For information about configuring the 7705 SAR for LSR ECMP, see the **lsr-load-balancing** and **system-ip-load-balancing** commands in the *7705 SAR Basic System Configuration Guide*, "System Information and General Commands" and the **lsr-load-balancing** command in the *7705 SAR Router Configuration Guide*, "Router Interface Commands".

For information about LDP tree-trace commands for tracing ECMP paths, see the *7705 SAR OAM and Diagnostics Guide*.



Note: LDP tree-trace works best with label-IP hashing (**lbl-ip**) enabled, instead of label-only (**lbl-only**) hashing. These options are set with the **lsr-load-balancing** command.



Note:

- Because of the built-in timeout to dynamic ARP, the MAC address of the remote peer needs to be renewed periodically. The flow of IP traffic resets the timers back to their maximum values. In the case of LDP ECMP, one link could be used for transporting user MPLS (pseudowire) traffic but the LDP session could possibly be using a different equal-cost link. For LDPs using ECMP and for static LSPs, it is important to ensure that the remote MAC address is learned and does not expire. Configuring static ARP entries or running continuous IP traffic ensures

that the remote MAC address is always known. Running BFD for fast detection of Layer 2 faults or running any OAM tools with SAA ensures that the learned MAC addresses do not expire.

- ARP entries are refreshed by static ARP and BFD, SAA, OSPF, IS-IS, or BGP.
- For information about configuring static ARP and running BFD, see the *7705 SAR Router Configuration Guide*.

6.1.9.1 Label operations

If an LSR is the ingress router for a specific IP prefix, LDP programs a PUSH operation for the prefix in the IOM. This creates an LSP ID to the Next Hop Label Forwarding Entry (NHLFE) mapping (LTN mapping) and an LDP tunnel entry in the forwarding plane. LDP also informs the tunnel table manager (TTM) about this tunnel. Both the LSP ID to NHLFE (LTN) entry and the tunnel entry have an NHLFE for the label mapping that the LSR received from each of its next-hop peers.

If the LSR is to function as a transit router for a specific IP prefix, LDP programs a SWAP operation for the prefix in the IOM. This involves creating an Incoming Label Map (ILM) entry in the forwarding plane. The ILM entry may need to map an incoming label to multiple NHLFEs.

If an LSR is an egress router for a specific IP prefix, LDP programs a POP entry in the IOM. This too results in an ILM entry being created in the forwarding plane, but with no NHLFEs.

When unlabeled packets arrive at the ingress LER, the forwarding plane consults the LTN entry and uses 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 a labeled packet arriving at a transit or egress LSR, the forwarding plane consults the ILM entry and either uses a hashing algorithm to map it to one of the NHLFEs if they exist (SWAP label) or routes the packet if there are no NHLFEs (POP label).

6.1.10 Graceful restart helper

Graceful restart is part of the LDP handshake process (that is, the LDP peering session initialization) and needs to be supported by both peers. Graceful restart helper provides a mechanism that allows the peers to cope with a service interruption because of a CSM switchover, which is a period of time when the standby CSM is not capable of synchronizing the states of the LDP sessions and labels being advertised and received.

Graceful restart helper decouples the data plane from the control plane so that if the control plane is not responding (that is, there is no LDP message exchange between peers), the data plane can still forward frames based on the last known (advertised) labels.

Because the 7705 SAR supports non-stop services / high-availability for LDP (and MPLS), the full implementation of graceful restart is not needed. However, graceful restart helper is implemented on the 7705 SAR to support non-high-availability devices. With graceful restart helper, if an LDP peer of the 7705 SAR requests graceful restart during the LDP handshake, the 7705 SAR agrees to it but does not request graceful restart. For the duration of the LDP session, if the 7705 SAR LDP peer fails, the 7705 SAR continues to forward MPLS packets based on the last advertised labels and does not declare the peer dead until the graceful restart timer expires.

6.1.11 Graceful handling of resource exhaustion

Graceful handling of resource exhaustion enhances the behavior of LDP when a data path or a CSM resource required for the resolution of a FEC is exhausted. In prior releases, the entire LDP protocol was shut down, causing all LDP peering sessions to be torn down and therefore impacting all peers. The user was required to fix the issue that caused the FEC scaling to be exceeded, and to restart the LDP session by executing the **no shutdown** CLI command. With graceful handling of resource exhaustion, only the responsible session or sessions are shut down, which impacts only the appropriate peer or peers.

Graceful handling of resources implements a capability by which the LDP interface to the peer, or the targeted peer in the case of a targeted LDP (T-LDP) session, is shut down.

If LDP tries to resolve a FEC over a link or a T-LDP session and runs out of data path or CSM resources, LDP brings down that interface or targeted peer, which brings down the Hello adjacency over that interface to all linked LDP peers or to the targeted peer. The interface is brought down for the LDP context only and is still available to other applications such as IP forwarding and RSVP LSP forwarding.

After taking action to free up resources, the user must manually perform a **no shutdown** command on the interface or the targeted peer to bring it back into operation. This re-establishes the Hello adjacency and resumes the resolution of FECs over the interface or to the targeted peer.

6.1.12 LDP support for unnumbered interfaces

Unnumbered interfaces are point-to-point interfaces that are not explicitly configured with a dedicated IP address and subnet; instead, they borrow (or link to) an IP address from another interface on the system (the system IP address, another loopback interface, or any other numbered interface) and use it as the source IP address for packets originating from the interface. For more information about support for unnumbered interfaces, see the *7705 SAR Router Configuration Guide*, "Unnumbered Interfaces".

This feature allows LDP to establish a Hello adjacency and to resolve unicast FECs over unnumbered LDP interfaces.

For example, LSR A and LSR B are the two endpoints of an unnumbered link. These interfaces are identified on each system with their unique link local identifier. The combination router ID and link local identifier uniquely identifies the interface in 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 that must originate from the interface. The borrowed IP address defaults to the system interface address, A and B in this example. The borrowed IP interface can be configured to any IP interface by using the following CLI command: **config> router>interface>unnumbered {ip-int-name | ip-address}**.

The **fec-originate** command, which defines how to originate a FEC for egress and non-egress LSRs, includes a parameter to specify the name of the interface that the label for the originated FEC is swapped to. For an unnumbered interface, this parameter is mandatory because an unnumbered interface does not have its own IP address.

When the unnumbered interface is added into LDP, the follow behavior occurs:

For link LDP (L-LDP) sessions:

1. The Hello adjacency is brought up using a link Hello packet with the source IP address set to the interface borrowed IP address and a destination IP address set to 224.0.0.2.
2. 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 are associated with a single LDP session.

3. The transport address for the TCP connection, which is encoded in the Hello packet, is always set to the LSR ID of the node whether the interface option was enabled using the **config>router>ldp>interface-parameters> interface>transport-address** command.
4. The **local-lsr-id** option can be configured on the interface and the value of the LSR ID can be changed to either the local interface or to some other interface name. If the local interface is selected or the provided interface name corresponds to an unnumbered IP interface, the unnumbered interface borrowed IP address is used as the LSR ID. In all cases, the transport address for the LDP session is updated to the new LSR ID value, but the link Hello packets continue to use the interface borrowed IP address as the source IP address.
5. The LSR with the highest transport address, the LSR ID in this case, bootstraps the TCP connection and LDP session.
6. The source and destination IP addresses of LDP packets are the transport addresses, that is, the LDP LSR IDs of the LSRs at the endpoints of the link (A and B in the example).

For targeted LDP (T-LDP) sessions:

1. The source and destination addresses of the targeted Hello packet are the LDP LSR IDs of systems A and B.
2. The **local-lsr-id** option can be configured on the interface for the targeted session and the value of the LSR ID can be changed to either the local interface or to some other interface name. If the local interface is selected or the provided interface name corresponds to an unnumbered IP interface, the unnumbered interface borrowed IP address is used as the LSR ID. In all cases, the transport address for the LDP session and the source IP address of the targeted Hello message are updated to the new LSR ID value.
3. The LSR with the highest transport address, the LSR ID in this case, bootstraps the TCP connection and LDP session.
4. The source and destination IP addresses of LDP packets are the transport addresses, that is, the LDP LSR IDs of the LSRs at the endpoints of the link (A and B in the example).

FEC resolution:

- LDP advertises/withdraws 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 the RTM cache. The next hop is 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 are 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 supported for numbered IP interfaces are supported for unnumbered interfaces, with the following exceptions:

- BFD is not supported on unnumbered IP interfaces
- LDP FRR is not triggered by a BFD session timeout, only by a physical failure or the local interface going down
- unnumbered IP interfaces cannot be added into LDP global and peer prefix policies

The unnumbered interface feature also extends the support of LSP ping and LSP traceroute to test an LDP unicast FEC that is resolved over an unnumbered LDP interface.

6.1.13 LDP FRR

LDP fast reroute (FRR) provides local protection for an LDP FEC by precalculating and downloading a primary and a backup NHLFE for the FEC to the LDP FIB. The primary NHLFE corresponds to the label of the FEC received from the primary next hop as per the standard LDP resolution of the FEC prefix in the RTM. The backup NHLFE corresponds to the label received for the same FEC from a loop-free alternate (LFA) next hop.

LDP FRR protects against single link or single node failure. SRLG failure protection is not supported.

Without FRR, when a local link or node fails, the router must signal the failure to its neighbors via the IGP providing the routing (OSPF or IS-IS), recalculate primary next-hop NHLFEs for all affected FECs, and update the FIB. Until the new primary next hops are installed in the FIB, any traffic destined for the affected FECs is discarded. This process can take hundreds of milliseconds.

LDP FRR improves convergence in case of a local link or node failure in the network, by using the label-FEC binding received from the LFA next hop to forward traffic for a specific prefix as soon as the primary next hop is not available. This means that a router resumes forwarding LDP packets to a destination prefix using the backup path without waiting for the routing convergence. Convergence times should be similar to RSVP-TE FRR, in the tens of milliseconds.

OSPF or IS-IS must perform the shortest path first (SPF) calculation of an LFA next hop, as well as the primary next hop, for all prefixes used by LDP to resolve FECs. The IGP also populates both routes in the RTM.

When LDP FRR is enabled and an LFA backup next hop exists for the FEC prefix in the RTM, or for the longest prefix the FEC prefix matches to when the **aggregate-prefix-match** option is enabled, LDP programs the data path with both a primary NHLFE and a backup NHLFE for each next hop of the FEC.

To perform a switchover to the backup NHLFE in the fast path, LDP follows the standard FRR failover procedures, which are also supported for RSVP-TE FRR.

When any of the following events occurs, the backup NHLFE is enabled for each affected FEC next hop:

- an LDP interface goes operationally down or is administratively shut down
In this case, LDP sends a neighbor/next hop down message to each LDP with which it has an adjacency over the interface.
- an LDP session to a peer goes down because the Hello timer or keepalive timer has expired over an interface
In this case, LDP sends a neighbor/next hop down message to the affected peer.
- the TCP connection used by a link LDP session to a peer goes down
In this case, LDP sends a neighbor/next hop down message to the affected peer.

See RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*, for more information about LFAs.

6.1.13.1 ECMP vs FRR

If ECMP is enabled, which provides multiple primary next hops for a prefix, LDP FRR is not used. That is, the LFA next hops are not populated in the RTM and the ECMP paths are used instead.

6.1.13.2 IGP shortcuts (RSVP-TE tunnels)

IGP shortcuts are an MPLS functionality where LSPs are treated like physical links within IGP; that is, LSPs can be used for next-hop reachability. If an RSVP-TE LSP is used as a shortcut by OSPF or IS-IS, it is included in the SPF calculation as a point-to-point link for both primary and LFA next hops. It can also be advertised to neighbors so that the neighboring nodes can also use the links to reach a destination via the advertised next hop.

IGP shortcuts can be used to simplify remote LFA support and simplify the number of LSPs required in a ring topology.

When both IGP shortcuts and LFA are enabled under OSPF or IS-IS, and LDP FRR is also enabled, the following applies:

- a FEC that is resolved to a direct primary next hop can be backed up by a tunneled LFA next hop
- a FEC that is resolved to a tunneled primary next hop does not have an LFA next hop; it relies on RSVP-TE FRR for protection

6.1.13.3 LDP FRR configuration

To configure LDP FRR, LFA calculation by the SPF algorithm must first be enabled under the OSPF or IS-IS protocol level with the command:

```
config>router>ospf>loopfree-alternates
```

or

```
config>router>ospf3>loopfree-alternates
```

or

```
config>router>isis>loopfree-alternates
```

Next, LDP must be enabled to use the LFA next hop with the command **config>router>ldp>fast-reroute**.

If IGP shortcuts are used, they must be enabled under the OSPF or IS-IS routing protocol. As well, they must be enabled under the MPLS LSP context, using the command **config>router>mpls>lsp>igp-shortcut**.

For information about LFA and IGP shortcut support for OSPF and IS-IS, see the *7705 SAR Routing Protocols Guide*, "LDP and IP fast reroute for OSPF prefixes" and "LDP and IP fast reroute for IS-IS prefixes".

Both LDP FRR and IP FRR are supported; for information about IP FRR, see the *7705 SAR Router Configuration Guide*, "IP FRR".

6.1.14 LDP-to-segment routing stitching for IPv4 /32 prefixes (IS-IS)

This feature provides stitching between an LDP FEC and an SR node SID route for the same IPv4 /32 IS-IS prefix by allowing the export of SR tunnels from the tunnel table manager (TTM) to LDP (IGP). In the LDP-to-SR data path direction, the LDP tunnel table route export policy supports the exporting of SR tunnels from the TTM to LDP.

A route policy option is configured to support LDP-to-SR stitching using the **config>router>policy-options** context. See the *7705 SAR Router Configuration Guide*, "Configuring LDP-to-Segment Routing Stitching".

Policies”, for a configuration example and to “Route Policy Command Reference” for information about the commands that are used.

After the route policy option is configured, the SR tunnels are exported from the TTM into LDP (IGP) using the **config>router>ldp>export-tunnel-table** command. See [LDP command reference](#) for more information about this command.

When configuring a route policy option, the user can restrict the exporting of SR tunnels from the TTM to LDP from a specific prefix list by excluding the prefix from the list.

The user can also restrict the exporting of SR tunnels from the TTM to a specific IS-IS IGP instance by specifying the instance ID in the **from protocol** statement. The **from protocol** statement is valid only when the protocol value is **isis**. Policy entries with any other protocol value are ignored when the route policy is applied. If the user configures multiple **from protocol** statements in the same policy or does not include the **from protocol** statement but adds a **default-action** of **accept**, then LDP routing uses the lowest instance ID in the IS-IS protocol to select the SR tunnel.

When the routing policy is enabled, LDP checks the SR tunnel entries in the TTM. Whenever an LDP FEC primary next hop cannot be resolved using an RTM route and an SR tunnel of type **isis** to the same destination IPv4 /32 prefix matches an entry in the export policy, LDP programs an LDP ILM and stitches it to the SR node SID tunnel endpoint. LDP then originates a FEC for the prefix and redistributes it to its LDP peer. When a LDP FEC is stitched to an SR tunnel, forwarded packets benefit from the protection of the LFA/remote LFA or TI-LFA backup next hop of the SR tunnel.

When resolving a FEC, LDP attempts a resolution in the RTM before attempting a resolution in the TTM, when both are available. That is, a swapping operation from the LDP ILM to an LDP NHLFE is attempted before stitching the LDP ILM to an SR tunnel endpoint.

In the SR-to-LDP data path direction, the SR mapping server provides a global policy for the prefixes corresponding to the LDP FECs the SR needs to stitch to. Therefore, a tunnel table export policy is not used. The user enables the exporting of the LDP tunnels for FEC prefixes advertised by the mapping server to an IGP instance using the command **config>router>isis>segment-routing>export-tunnel-table ldp**. See the *7705 SAR Routing Protocols Guide*, “IS-IS Command Reference”, for more information about this command.

When the **export-tunnel-table ldp** command is enabled, the IGP monitors the LDP tunnel entries in the TTM. Whenever an IPv4 /32 LDP tunnel destination matches a prefix for which the IGP received a prefix SID sub-TLV from the mapping server, the IGP instructs the SR module to program the SR ILM and to stitch it to the LDP tunnel endpoint. The SR ILM can stitch to an LDP FEC resolved over the LDP link. When an SR tunnel is stitched to an LDP FEC, forwarded packets benefit from the protection of the LFA backup next hop of the LDP FEC.

When resolving a node SID, the IGP attempts a resolution of the prefix SID received in an IP reachability TLV before attempting a resolution of a prefix SID received via the mapping server, when both are available. That is, a swapping operation of the SR ILM to an SR NHLFE is attempted before stitching it to an LDP tunnel endpoint. See the *7705 SAR Routing Protocols Guide*, “Prefix SID Resolution for a Segment Routing Mapping Server”, for more information about prefix SID resolution.


It is recommended that the **bfd-enable** option be enabled on the interfaces for both LDP and IGP contexts to speed up the failure detection and the activation of the LFA/remote LFA backup next hop in either direction. This applies particularly for remote failures. For the LDP context, the **config>router>ldp>interface-parameters>interface>bfd-enable** command string is used; see [LDP commands](#). For the IGP context, the **config>router>isis>interface>bfd-enable** command string is used; see the *7705 SAR Routing Protocols Guide*, “IS-IS Command Reference”.

The sections that follow describe how stitching is performed in the LDP-to-SR and SR-to-LDP data path directions.

6.1.14.1 Stitching in the LDP-to-SR direction

Stitching in the data plane in the LDP-to-SR direction is based on the LDP module monitoring the TTM for an SR tunnel of a prefix matching an entry in the LDP TTM export policy.

Figure 44: Stitching in the LDP-to-SR direction



In the figure, router R1 is at the boundary between an SR domain and an LDP domain and is configured to stitch between SR and LDP. Link R1-R2 is LDP-enabled, but router R2 does not support SR or SR is disabled.

The following steps are performed by the boundary router R1 to configure stitching:

1. Router R1 receives a prefix SID sub-TLV in an IS-IS IP reachability TLV originated by router Ry for prefix Y.
2. R1 resolves the prefix SID and programs an NHLFE on the link toward the next hop in the SR domain. R1 programs an SR ILM and points it to the NHLFE.
3. Because R1 is programmed to stitch LDP to SR, LDP in R1 checks the TTM and finds the SR tunnel to prefix Y. LDP programs an LDP ILM and points it to the SR tunnel. As a result, both the SR ILM and LDP ILM are now pointing to the SR tunnel, one via the SR NHLFE and the other via the SR tunnel endpoint.
4. R1 advertises the LDP FEC for prefix Y to all its LDP peers. R2 is now able to install an LDP tunnel toward Ry.
5. If R1 finds multiple SR tunnels to destination prefix Y, R1 uses the lowest instance ID in the IS-IS protocol to select the tunnel.
6. If the user configured multiple **from** statements or did not include the **from** statement but added a default action of **accept** for the IS-IS protocol, R1 selects the tunnel to destination prefix Y by using the lowest instance ID in the IS-IS protocol.



Note: If R1 has already resolved an LDP FEC for prefix Y, it has an ILM assigned to it. However, this ILM is not updated to point toward the SR tunnel because LDP attempts a resolution in the RTM before attempting a resolution in the TTM. Therefore, an LDP tunnel is selected before an SR tunnel. Similarly, if an LDP FEC is received after the stitching is

programmed, the LDP ILM is updated to point to the LDP NHLFE because LDP is able to resolve the LDP FEC in the RTM.

7. The user enables SR in R2. R2 resolves the prefix SID for prefix Y and installs the SR ILM and the SR NHLFE. R2 is now able to forward packets over the SR tunnel to router Ry. There is no activity in R1 because the SR ILM is already programmed.
8. The user disables LDP over the R1-R2 interface in both directions. This causes the LDP FEC ILM and NHLFE to be removed in R1 and in R2, which can then only do forwarding using the SR tunnel toward Ry.

6.1.14.2 Stitching in the SR-to-LDP direction

Stitching in the data plane in the SR-to-LDP direction is based on the IGP monitoring the TTM for an LDP tunnel of a prefix matching an entry in the SR TTM export policy.

In [Figure 44: Stitching in the LDP-to-SR direction](#), router R1 is at the boundary between an SR domain and an LDP domain and is configured to stitch between SR and LDP. Link R1-R2 is LDP-enabled but router R2 does not support SR or SR is disabled.

The following steps are performed by the boundary router R1 to configure stitching:

1. R1 receives an LDP FEC for prefix X from router Rx in the LDP domain. The RTM in R1 indicates that the interface to R2 is the next hop for prefix X.
2. LDP in R1 resolves the received FEC in the RTM and creates an LDP ILM for the FEC with an ingress label (for example, label L1), and points it to an LDP NHLFE toward R2 with egress label L2.
3. R1 receives a prefix SID sub-TLV from the R5 mapping server for prefix X.
4. The IGP in R1 attempts to resolve in its routing table the next hop of prefix X over the interface to R2. R1 detects that R2 did not advertise support of SR and therefore the SID resolution for prefix X in the routing table fails.
5. The IGP in R1 then attempts to resolve the prefix SID of prefix X in the TTM because it detects that it is configured for SR-to-LDP stitching. R1 finds an LDP tunnel to prefix X in the TTM, instructs the SR module to program an SR ILM with ingress label L3, and points it to the LDP tunnel endpoint, therefore stitching ingress label L3 to egress label L2.



Note:

- The ILMs for LDP and SR are both pointing to the same LDP tunnel, one via NHLFE and one via the tunnel endpoint.
- No SR tunnel to destination prefix X should be programmed in the TTM following the resolution of the prefix SID of prefix X in the TTM.
- If the IGP is not able to resolve the SID resolution for prefix X in [step 4](#) and [step 5](#), a trap is generated for the prefix SID resolution failure. An existing trap for the prefix SID resolution failure is enhanced to state whether the prefix SID that failed the resolution attempts was part of a mapping server TLV or an IP reachability TLV.

6. The user enables segment routing on R2.
7. The IGP in R1 discovers that R2 supports SR.

Because R1 still has a prefix SID for prefix X from the mapping server R5, it maintains the stitching of the SR ILM for prefix X to the LDP FEC.

8. The user disables the LDP interface between R1 and R2 in both directions. This causes the LDP FEC ILM and NHLFE for prefix X to be removed in R1 and triggers the re-evaluation of the SIDs.
9. R1 first attempts the resolution in the routing table. Because the next hop for prefix X supports SR, the IGP instructs the SR module to program an NHLFE for the prefix SID of prefix X with egress label L4 and with an outgoing interface to R2. R1 creates an SR tunnel in the TTM for destination prefix X. R1 also changes the SR ILM with ingress label L3 to point to the SR NHLFE with egress label L4.
Router R2 now becomes the SR-LDP stitching router.
10. Router Rx, which owns prefix X, is upgraded to support SR. Rx sends a prefix SID sub-TLV to R1 in an IS-IS IP reachability TLV for prefix X. The SID information may or may not be the same as the information received from the mapping server R5. If the SID information is not the same, the IGP in R1 chooses the prefix SID originated by Rx and updates the SR ILM and NHLFE with the appropriate labels.
11. The user then cleans up the mapping server and removes the mapping entry for prefix X, which is then withdrawn by IS-IS.

6.1.14.3 TTL propagation and ICMP tunneling

When stitching is performed between an LDP FEC and an SR IS-IS node SID tunnel, the TTL of the outer LDP or SR label is decreased, similar to a regular swapping operation at an LSR.

6.1.15 LDP FRR remote LFA and TI-LFA backup using an SR tunnel for IPv4 /32 prefixes (IS-IS)

This feature allows an SR tunnel to be used as a remote LFA or TI-LFA backup tunnel next hop by an LDP FEC. The feature is enabled using the CLI command string **config>router>ldp>fast-reroute backup-sr-tunnel**. See [LDP commands](#) for more information.

This feature requires the [LDP-to-segment routing stitching for IPv4 /32 prefixes \(IS-IS\)](#) feature as a prerequisite, because the LSR performs the stitching of the LDP ILM to an SR tunnel when the primary LDP next hop of the FEC fails. Therefore, LDP monitors SR tunnels programmed by the IGP in the TTM without the need for a mapping server.

It is assumed that:

- the **backup-sr-tunnel** option is enabled in LDP
- the **loopfree-alternates ti-lfa** or **loopfree-alternates remote-lfa** option is enabled in the IGP instance (see the *7705 SAR Routing Protocols Guide*, "IS-IS Command Reference")



Note: The **loopfree-alternates** options can be enabled separately or together. If both options are enabled, TI-LFA backup takes precedence over remote LFA backup.

- LDP was able to resolve the primary next hop of the LDP FEC in the RTM

If the IGP LFA SPF does not find a regular LFA backup next hop for an LDP FEC prefix, it runs the TI-LFA and remote LFA algorithms. If the IGP LFA SPF finds a remote LFA or TI-LFA tunnel next hop, LDP programs the primary next hop of the FEC using an LDP NHLFE and programs the remote LFA or TI-LFA backup tunnel next hop using an LDP NHLFE pointing to the SR tunnel endpoint.



Note: The LDP packet is not sent over the SR tunnel. The LDP label is stitched to the segment routing label stack. LDP points both the LDP ILM and the LTN to the backup LDP NHLFE, which uses the SR tunnel endpoint.

6.1.15.1 Feature behavior

- When LDP resolves a primary next hop in the RTM or a remote LFA or TI-LFA backup next hop using an SR tunnel in the TTM, LDP programs a primary LDP NHLFE and a backup LDP NHLFE with an implicit null label value pointing to the SR tunnel that has the remote LFA or TI-LFA backup programmed for the same prefix.
- If the LDP FEC primary next hop fails and LDP has preprogrammed a remote LFA and TI-LFA next hop with an LDP backup NHLFE pointing to an SR tunnel, the LDP ILM/LTN switches to it.



Note: If the LDP FEC primary next hop failure impacts only the LDP tunnel primary next hop but not the SR tunnel primary next hop, the LDP backup NHLFE points to the primary next hop of the SR tunnel; the LDP ILM/LTN traffic follows this path instead of the remote LFA or TI-LFA next hop of the SR tunnel until the remote LFA or TI-LFA next hop is activated.

- If the LDP FEC primary next hop becomes unresolved in the RTM, LDP switches the resolution to an SR tunnel in the TTM, if one exists, following the steps described in [Stitching in the LDP-to-SR direction](#).
- If both the LDP primary next hop and a regular LFA next hop become resolved in RTM, the LDP FEC programs the primary NHLFE and backup NHLFE.

6.1.16 TCP MD5 and enhanced TCP authentication

The operation of a network can be compromised if an unauthorized system is able to form or hijack an LDP session and inject control packets by falsely representing itself as a valid neighbor. This risk can be mitigated by enabling TCP MD5 authentication, the TCP Enhanced Authentication Option, or the TCP Authentication Option (TCP-AO) on one or more of the sessions. Only one authentication type can be configured at a time for a TCP connection.

When TCP MD5 authentication is enabled on a session, every TCP segment exchanged with the peer includes a TCP option (19) containing a 16-byte MD5 digest of the segment (more specifically the TCP/IP pseudo-header, TCP header, and TCP data). The MD5 digest is generated and validated using an authentication key that must be known to both sides. If the received digest value is different from the locally computed one, the TCP segment is dropped, thereby protecting the router from a spoofed TCP segment.

The TCP Enhanced Authentication Option, as defined in *draft-bonica-tcpauth-05.txt*, and TCP-AO, as defined in RFC 5925, are TCP extensions that enhance security for BGP and LDP, improving upon the authentication strategy offered by TCP MD5. These extensions allow TCP peers to authenticate messages passed between one another. They include the ability to change keys in a BGP or LDP session seamlessly without tearing down the session and allow for stronger authentication algorithms to be used to authenticate routing messages. They are intended for applications where secure administrative access to both endpoints of the TCP connection is normally available.

TCP-AO and the TCP Enhanced Authentication Option use keychains that are associated with every protected TCP connection. Keychains are configured in the **config>system>security>keychain** context. For more information about configuring keychains, see the *7705 SAR System Management Guide*, "TCP enhanced authentication and keychain authentication".

6.2 LDP point-to-multipoint support

The 7705 SAR supports point-to-multipoint LDP. This section contains information about the following topics:

- [LDP point-to-multipoint configuration](#)
- [LDP point-to-multipoint protocol](#)
- [Make-before-break \(MBB\)](#)
- [ECMP support](#)
- [Inter-AS non-segmented mLDP](#)
- [Intra-AS non-segmented mLDP](#)

6.2.1 LDP point-to-multipoint configuration

A node running LDP also supports point-to-multipoint LSP setup using LDP. By default, the node advertises the capability to a peer node using the point-to-multipoint capability TLV in LDP initialization message.

The **multicast-traffic** configuration option (per interface) restricts or allows the use of an interface for LDP multicast traffic forwarding toward a downstream node. The **interface** configuration option does not restrict or allow the exchange of the point-to-multipoint FEC by way of an established session to the peer on an interface, but only restricts or allows the use of next hops over the interface.

6.2.2 LDP point-to-multipoint protocol

Only a single generic identifier range is defined for signaling a multipoint data tree (MDT) for all client applications. Implementation on the 7705 SAR reserves the range 1 to 8292 for generic point-to-multipoint LSP ID values for static point-to-multipoint LSP on the root node.

6.2.3 Make-before-break (MBB)

After a transit or leaf node detects that the upstream node toward the root node of a multicast tree has changed, the node follows the graceful procedure that allows make-before-break transition to the new upstream node. Make-before-break support is optional via the **mp-mbb-time** command. If the new upstream node does not support MBB procedures, the downstream node waits for the configured timer to time out before switching over to the new upstream node.

6.2.4 ECMP support

If multiple ECMP paths exist between two adjacent nodes, then the upstream node of the multicast receiver programs all entries in the forwarding plane. Only one entry is active and it is based on the ECMP hashing algorithm.

6.2.5 Inter-AS non-segmented mLDP

This feature allows multicast services to use segmented protocols and span them over multiple autonomous systems (ASs) in the same way as unicast services. Because IP VPN or GRT services span multiple IGP areas or multiple ASs, either for a network designed to deal with scale or as result of commercial acquisitions, operators may require inter-AS VPN (unicast) connectivity. For example, an inter-AS VPN can break the IGP, MPLS, and BGP protocols into access segments and core segments, allowing higher scaling of protocols by segmenting them into their own islands. The 7705 SAR allows for a similar provisioning of multicast services and for spanning these services over multiple IGP areas or multiple ASs.

Multicast LDP (mLDP) supports non-segmented mLDP trees for inter-AS solutions that are applicable for NG-MVPN services.


6.2.5.1 LDP recursive FEC process

For inter-AS networks where the leaf node does not have the root in the RTM or where the leaf node has the root in the RTM using BGP, and the leaf's local AS intermediate nodes do not have the root in their RTM because they are not BGP-enabled, RFC 6512 defines a recursive opaque value and procedure for LDP to build an LSP through multiple ASs.

For mLDP to be able to signal through a multiple-AS network where the intermediate nodes do not have a routing path to the root, a recursive opaque value is needed. The LDP FEC root resolves the local ASBR, and the recursive opaque value contains the point-to-multipoint FEC element, encoded as specified in RFC 6513, with a type field, a length field, and a value field of its own.


RFC 6826 section 3 defines the Transit IPv4 opaque for point-to-multipoint LDP FEC, where the leaf in the local AS wants to establish an LSP to the root for point-to-multipoint LSP. The following figure shows FEC in this network configuration.

Figure 45: mLDP FEC for single AS with transit IPv4 opaque



The following figure shows an inter-AS FEC with recursive opaque based on RFC 6512.

Figure 46: MLDP FEC for inter-AS with recursive opaque value




sw1302

As shown in the figure, the root "10.0.0.21" is an ASBR and the opaque value contains the original mLDP FEC. Therefore, in the leaf AS where the actual root "10.0.0.14" is not known, the LDP FEC can be routed using the local root of ASBR. When the FEC arrives at an ASBR that is colocated in the same AS as the actual root, an LDP FEC with transit IPv4 opaque is generated.

The following figure shows an end-to-end example of a non-VPN mLDP with recursive opaque for inter-AS.

Figure 47: Non-VPN mLDP with recursive opaque for inter-AS



sw1299

As shown in the figure, the leaf is in AS3s where the AS3 intermediate nodes do not have ROOT-1 in their RTM. The leaf has S1 installed in the RTM via BGP. All ASBRs are acting as next-hop-self in the BGP domain. The leaf resolving S1 via BGP generates an mLDP FEC with recursive opaque, represented as:

Leaf FEC: <Root=ASBR-3, opaque-value=<Root=Root-1, <opaque-value = S1,G1>>>

This FEC is routed through the AS3 Core to ASBR-3.



Note: AS3 intermediate nodes are not BGP-capable because they do not have ROOT-1 in their RTM.

At ASBR-3, the FEC is changed to:

ASBR-3 FEC: <Root=ASBR-1, opaque-value=<Root=Root-1, <opaque-value = S1,G1>>>

This FEC is routed from ASBR-3 to ASBR-1. ASBR-1 is colocated in the same AS as ROOT-1. Therefore, ASBR-1 does not need a FEC with a recursive opaque value.

ASBR-1 FEC: <Root=Root-1, <opaque-value =S1,G1>>

This process allows all multicast services to work over inter-AS networks.

6.2.5.1.1 Supported recursive opaque values

A recursive FEC is built using the recursive opaque value (opaque value 7). All non-recursive opaque values can be recursively embedded into a recursive opaque.

The following table lists the supported opaque values for the 7705 SAR.


Table 43: Opaque types

Opaque type	Opaque name	RFC	7705 SAR use	FEC representation
1	Generic LSP identifier	RFC 6388	VPRN Local AS	<Root, Opaque<P2MPID>>
7	Recursive opaque value	RFC 6512	Inter-AS option C MVPN over mLDP	<ASBR, Opaque<Root, Opaque<P2MPID>>>

6.2.5.1.2 Redundancy and resiliency

For mLDP, MoFRR is only supported within a single AS. In an inter-area option C scenario, MoFRR is not supported end-to-end with PEs crossing ASBR boundaries, as shown in the following figure.

Figure 48: ASBRs using EBGP without IGP




Non-segmented mLDP functions with ASBRs directly connected or connected via an IGP domain, as shown in the figure.

6.2.5.2 OAM

LSPs are unidirectional tunnels. When an LSP ping is sent, the echo request is transmitted via the tunnel and the echo response is transmitted as an IP packet to the source. Similarly, for a **p2mp-lsp-ping**, on the root, the echo request is transmitted via the mLDP point-to-multipoint tunnel to all leafs and the leafs use the IP packet to respond to the root.

The echo request for mLDP is generated carrying a root Target FEC Stack TLV (see the following figure), which is used to identify the multicast point-to-multipoint LSP under test at the leaf. The Target FEC Stack TLV must carry an mLDP point-to-multipoint FEC Stack Sub-TLV from RFC 6388 or RFC 6512.

Figure 49: ECHO request target FEC stack TLV



The same concept applies to inter-AS and non-segmented mLDP. The leafs in the remote AS should be able to resolve the root via GRT routing. This is possible for inter-AS option C where the root is usually in the leaf RTM, which is a next-hop ASBR.

OAM functionality for option C is summarized in the following table.


Table 44: OAM functionality for option C

OAM command (for mLDP)	Leaf and root in same AS	Leaf and root in different AS (option C)
p2mp-lsp-ping ldp	✓	✓

6.2.5.3 ECMP support

In the following figure, the leaf discovers ROOT-1 from all three ASBRs (ASBR-3, ASBR-4 and ASBR-5).

Figure 50: ECMP support



The leaf chooses which ASBR is used for the multicast stream using the following process:

1. The leaf determines the number of ASBRs that should be part of the hash calculation.

The number of ASBRs that are part of the hash calculation comes from the ECMP value (configured with the **config>router>ecmp** command). For example, if the ECMP value is set to 2, only two of the ASBRs are part of the hash algorithm selection.

2. After selecting the upstream ASBR, the leaf determines whether there are multiple equal cost paths between it and the chosen ASBR.

- If there are multiple ECMP paths between the leaf and the ASBR, the leaf performs another ECMP selection based on the configured ECMP value. This is a recursive ECMP lookup.
- The first lookup chooses the ASBR and the second lookup chooses the path to that ASBR.

For example, if ASBR 5 in the figure was chosen, there are three paths between the leaf and ASBR-5. A second ECMP decision is therefore made to choose the path.

3. At ASBR-5, the process is repeated. For example, in the figure, ASBR-5 goes through steps 1 and 2 to choose between ASBR-1 and ASBR-2 and perform a second recursive ECMP lookup to choose the path to that ASBR.

When there are several candidate upstream LSRs, the LSR must select one upstream LSR. The algorithm used for the LSR selection is a local decision. If the LSR selection is done over a LAN interface, the procedure described in [ECMP hash algorithm](#) should be applied to ensure that the same upstream LSR is elected from the set of candidate receivers on that LAN.

The ECMP hash algorithm ensures that there is a single forwarder over the LAN for a particular LSP.

6.2.5.3.1 ECMP hash algorithm

The ECMP hash algorithm requires the opaque value of the FEC (see [Table 43: Opaque types](#)) and is based on RFC 6388 section 2.4.1.1.


The following hash is performed: $H = (\text{CRC32}(\text{Opaque Value})) \bmod N$, where N is the number of upstream LSRs. The Opaque Value is the field identified in the FEC element after Opaque Length. The Opaque Length field indicates the size of the opaque value used in this calculation.

The candidate upstream LSRs are numbered from lowest to highest IP addresses. The selected upstream LSR is the LSR that has the value H , as calculated by the hash algorithm.

6.2.6 Intra-AS non-segmented mLDP

Non-segmented mLDP intra-AS (inter-area) is supported on option C only. The following figure shows a typical intra-AS topology with a backbone IGP area 0 and access non-backbone IGP areas 1 and 2. In these topologies, the ABRs usually do a next-hop-self for BGP label routes, which requires recursive FEC.

Figure 51: Intra-AS non-segmented topology



6.2.6.1 ABR MoFRR for intra-AS

With ABR multicast-only fast reroute (MoFRR) in the intra-AS environment, the leaf chooses a local primary ABR and a backup ABR, with separate mLDP signaling toward these two ABRs. In addition, each path from a leaf to the primary ABR and from a leaf to the backup ABR supports IGP MoFRR.

End-to-end intra-AS MoFRR is not supported and stops at each IGP instance. MoFRR stops at the ABR if the ABR has a leaf.

6.2.6.2 Interaction with an inter-AS non-segmented mLDP solution

Intra-AS option C is supported in conjunction with inter-AS option C. Intra-AS option C with inter-AS option B is not supported.

6.3 MLDP fast upstream switchover

This feature allows a downstream LSR of an mLDP FEC to perform a fast switchover to source the traffic from another upstream LSR while IGP and LDP are converging because of a failure of the upstream LSR, where the upstream LSR is the primary next hop of the root LSR for the point-to-multipoint FEC. The feature is enabled using the **mcast-upstream-frr** command.

The feature provides upstream fast reroute (FRR) node protection for mLDP FEC packets. The protection is at the expense of traffic duplication from two different upstream nodes into the node that performs the fast upstream switchover.

The detailed procedures for this feature are described in *draft-pdutta-mpls-mldp-up-redundancy*.

6.3.1 MLDP fast upstream switchover configuration

To enable the mLDP fast upstream switchover feature, configure the following option in the CLI:

```
config>router>ldp>mcast-upstream-frr
```

When **mcast-upstream-frr** is enabled and LDP is resolving an mLDP FEC received from a downstream LSR, LDP checks for the existence of an ECMP next hop or a loop-free alternate (LFA) next hop to the root LSR node. If LDP finds one, it programs a primary incoming label map (ILM) on the interface corresponding to the primary next hop and a backup ILM on the interface corresponding to the ECMP or LFA next hop. LDP then sends the corresponding labels to both upstream LSR nodes. In normal operation, the primary ILM accepts packets and the backup ILM drops them. If the interface or the upstream LSR of the primary ILM goes down, causing the LDP session to go down, the backup ILM starts accepting packets.

To use the ECMP next hop, configure the **ecmp max-ecmp-routes** value in the system to be at least 2, using the following command:

```
config>router>ecmp max-ecmp-routes
```

To use the LFA next hop, enable LFA using the following commands (as needed):

```
config>router>isis>loopfree-alternates
```

or

```
config>router>ospf>loopfree-alternates
```

Enabling IP FRR or LDP FRR is not strictly required, because LDP only needs to know the location of the alternate next hop to the root LSR to send the label mapping message and program the backup ILM during the initial signaling of the tree. That is, enabling the LFA option is sufficient for providing the backup ILM information. However, if unicast IP and LDP prefixes need to be protected, then IP FRR and LDP FRR—and the mLDP fast upstream switchover—can be enabled concurrently using the following commands:

```
config>router>ip-fast-reroute
```

or

```
config>router>ldp>fast-reroute
```

An mLDP FRR fast switchover relies on the fast detection of a lost LDP session to the upstream peer to which the primary ILM label had been advertised. To ensure fast detection of a lost LDP session, do the following:

1. Enable BFD on all LDP interfaces to upstream LSR nodes. When BFD detects the loss of the last adjacency to the upstream LSR, it brings down the LDP session immediately, which causes the CSM to activate the backup ILM.
2. If there is a concurrent T-LDP adjacency to the same upstream LSR node, enable BFD on the T-LDP peer in addition to enabling it on the interface.
3. Enable the **ldp-sync-timer** option on all interfaces to the upstream LSR nodes. If an LDP session to the upstream LSR to which the primary ILM is resolved goes down for any reason other than a failure of the interface or the upstream LSR, then routing and LDP go out of synchronization. This means that the backup ILM remains activated until the next time SPF is run by IGP.

By enabling the IGP-LDP synchronization feature, the advertised link metric changes to the maximum value as soon as the LDP session goes down. This, in turn, triggers an SPF, and LDP likely downloads a new set of primary and backup ILMs.

6.3.2 MLDP fast upstream switchover behavior


This feature allows a downstream LSR to send a label binding to two upstream LSR nodes but only accept traffic as follows:

- for normal operation, traffic is accepted from the ILM on the interface to the primary next hop of the root LSR for the point-to-multipoint FEC
- for failure operation, traffic is accepted from the ILM on the interface to the backup next hop

A candidate upstream LSR node must be either an ECMP next hop or an LFA next hop. Either option allows the downstream LSR to perform a fast switchover and to source the traffic from another upstream LSR while IGP is converging because of a failure of the LDP session of the upstream peer, which is the primary next hop of the root LSR for the point-to-multipoint FEC. That is, the candidate upstream LSR node provides upstream FRR node protection for the mLDP FEC packets.

MLDP fast upstream switchover is illustrated in the following figure. LSR U is the primary next hop for the root LSR R of the point-to-multipoint FEC. LSR U' is an ECMP or LFA backup next hop for the root LSR R of the same point-to-multipoint FEC.

Figure 52: MLDP fast upstream switchover



25937

In the figure, downstream LSR Z sends a label mapping message to both upstream LSR nodes, and programs the primary ILM on the interface to LSR U and the backup ILM on the interface to LSR U'. The labels for the primary and backup ILMs must be different. Thus LSR Z attracts traffic from both ILMs. However, LSR Z blocks the ILM on the interface to LSR U' and only accepts traffic from the ILM on the interface to LSR U.

If the link to LSR U fails, or LSR U fails, causing the LDP session to LSR U to go down, LSR Z detects the failure and reverse the ILM blocking state. In addition, LSR Z immediately starts receiving traffic from LSR U' until IGP converges and provides a new primary next hop and a new ECMP or LFA backup next hop, which may or may not be on the interface to LSR U'. When IGP convergence is complete, LSR Z updates the primary and backup ILMs in the datapath.



Note: LDP uses the interface of either an ECMP next hop or an LFA next hop to the root LSR prefix, whichever is available, to program the backup ILM. However, ECMP next hop and LFA next hop are mutually exclusive for a specific prefix. IGP installs the ECMP next hop in preference to the LFA next hop as a prefix in the routing table manager (RTM).

If one or more ECMP next hops for the root LSR prefix exist, LDP picks the interface for the primary ILM based on the rules of mLDP FEC resolution specified in RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*:

1. The candidate upstream LSRs are numbered from lowest to highest IP address.
2. The following hash is performed:

$$H = (\text{CRC32}(\text{Opaque Value})) \bmod N$$

where **N** is the number of upstream LSRs

The Opaque Value is the field in the point-to-multipoint FEC element immediately after the Opaque Length field. The Opaque Length indicates the opaque value used in this calculation.

3. The selected upstream LSR U is the LSR that has the number **H**.

LDP then picks the interface for the backup ILM using the following new rules:

```
if (H + 1 < NUM_ECMP) {
```

```
// If the hashed entry is not last in the next hops then pick up the next as backup.
```

```
  backup = H + 1;
```

```
} else {
```

```
// Wrap around and pick up the first.
```

```
  backup = 1;
```

```
}
```


In some topologies, it is possible that no ECMP or LFA next hop is found. In this case, LDP programs the primary ILM only.

6.4 LDP IPv6

The 7705 SAR extends the LDP control plane and data plane to support LDP IPv6 adjacencies and sessions using 128-bit LSR ID.

The implementation allows for concurrent support of independent LDP IPv4 (which uses a 32-bit LSR ID) and LDP IPv6 adjacencies and sessions between peer LSRs and over the same interfaces or different set of interfaces.

The following figure shows an example of an LDP adjacency and session over an IPv6 interface.

Figure 53: LDP adjacency and session over an IPv6 interface

27908

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, LSR-A/128 and LSR-B/128 use the system interface IPv6 address.

Although the LDP control plane can operate using only the IPv6 system address, it is recommended that the user configure the IPv4-formatted router ID in order for OSPF, IS-IS, and BGP to operate properly.

The following sections describe LDP IPv6 behavior on the 7705 SAR:

- [Link LDP](#)
- [Targeted LDP](#)
- [FEC resolution](#)
- [LDP session capabilities](#)
- [LDP adjacency capabilities](#)
- [IP address and FEC distribution](#)
- [IGP and static route synchronization with LDP](#)
- [BFD operation](#)
- [Services using SDP with an LDP IPv6 FEC](#)
- [Mirror services](#)
- [OAM support with LDP IPv6](#)
- [Interoperability](#)
- [Upgrading from IPv4 to IPv6](#)

6.4.1 Link LDP

LDP IPv6 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 a 32-bit LSR ID, as defined in draft-ietf-mpls-ldp-ipv6-14.

A Hello adjacency is brought up using a link Hello packet with a 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 by default to the LSR ID of the LSR. The transport address is instead set to the interface IPv6 address if the user enables the **interface** option in one of the following contexts:

- **config>router>ldp>if-params>ipv6>transport-address**
- **config>router>ldp>if-params>if>ipv6>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 another interface name, including loopback. The global unicast IPv6 address corresponding to the primary IPv6 address of the interface is used as the LSR ID. If the interface 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.

The source and destination addresses of LDP/TCP session packets are the IPv6 transport addresses.

6.4.2 Targeted LDP

The source and destination addresses of targeted Hello packets are the LDP IPv6 LSR- IDs of systems A and B in [Figure 53: 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, including loopback or not. 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 targeted Hello messages are updated with 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.

The source and destination IP addresses of LDP/TCP session packets are the IPv6 transport addresses the LDP LSR IDs of systems A and B in [Figure 53: LDP adjacency and session over an IPv6 interface](#).

6.4.3 FEC resolution

LDP advertises and withdraws all interface IPv6 addresses using the Address/ Address-Withdraw message. Both the link local unicast address and the configured global unicast addresses of an interface are advertised.

Like LDP IPv4 sessions, LDP FEC types can be exchanged over an LDP IPv6 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 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.

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 or spoke SDP auto-creation from a service such as BGP-AD VPLS, BGP-VPWS, or dynamic MS- PW.

6.4.4 LDP session capabilities

LDP can advertise all FEC types over an LDP IPv4 or an LDP IPv6 session. The FEC types are: IPv4 prefix FEC, IPv6 prefix FEC, IPv4 P2MP FEC (with MVPN), and PW FEC 128.

LDP also supports signaling the enabling or disabling of the advertisement of the following subset of FEC types during the LDP IPv4 or IPv6 session initialization phase, and when the session is already up:

- **IPv4 prefix FEC**

This is performed using the State Advertisement Control (SAC) capability TLV as specified in draft-ietf-mpls-ldp-ip-pw-capability. The SAC capability TLV includes the IPv4 SAC element having the D-bit (Disable-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 State Advertisement Control (SAC) capability TLV as specified in draft-ietf-mpls-ldp-ip-pw-capability. The SAC capability TLV includes the IPv6 SAC element having the D-bit (Disable-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 to update the state of this FEC type.

- **P2MP FEC (IPv4 only)**

This is performed using the P2MP capability TLV as specified in RFC 6388. The P2MP capability TLV has the S-bit (State-bit) with a value of set or reset to enable or disable this FEC type respectively.

The LSR can send this TLV in the LDP initialization message and, subsequently, in an LDP capability message to update the state of this FEC type.

During LDP session initialization, each LSR indicates to its peers which FEC type it supports by including the capability TLV for it in the LDP initialization message. The 7705 SAR enables the IPv4 and IPv6 Prefix FEC types by default and sends their 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 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.

Dynamic Capability, as defined in RFC 5561, allows all FEC types to update the enabled or disabled state after the LDP session initialization phase. An LSR informs its peer that it supports Dynamic Capability by including the Dynamic Capability Announcement TLV in the LDP initialization message. If both LSRs advertise this capability, the user can enable or disable any of the above 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 (Disable-bit) set or reset, or the P2MP capability TLV (IPv4 only) in a capability message with the S-bit (State-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 a specific LDP session to a peer by using the following CLI commands:

- **config>router>ldp>session-params>peer>fec-type-capability>prefix-ipv4**
- **config>router>ldp>session-params>peer>fec-type-capability>prefix-ipv6**
- **config>router>ldp>session-params>peer>fec-type-capability>p2mp**

6.4.5 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 specified 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, then FECs are not advertised and enabling that FEC type at the adjacency level do 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 (Disable-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 specified interface to a peer. This provides the ability to dedicate links and data path resources to specific FEC types. For IPv4 and IPv6 prefix FECs, a subset of ECMP links to an LSR peer may be configured to carry one of the two FEC types. An mLDP P2MP FEC (IPv4 only) 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 the disabling of a capability in the LDP Hello message, no FECs of the corresponding FEC type is resolved by either peer over the link of this adjacency for the lifetime of the LDP Hello adjacency, unless one or both peers sends the LDP adjacency capability TLV subsequently to explicitly enable it.

The user can enable or disable a FEC type for a specific LDP interface to a peer by using the following CLI commands:

- **config>router>ldp>if-params>if>ipv4>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 to the IPv4 P2MP FEC, deprecate the existing **multicast-traffic** command under the interface. Unlike the session-level capability, these commands can disable multicast FEC for IPv4.

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
- 4 : P2MP IPv6 forwarding (not supported on the 7705 SAR)
- 5 : MP2MP IPv4 forwarding
- 6 : MP2MP IPv6 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.

6.4.6 IP address and FEC distribution

When an LDP LSR initializes the LDP session to the peer LSR and the session comes up, IP addresses and FECs are distributed. Local IPv4 and IPv6 interface addresses are exchanged using the Address and Address Withdraw messages. 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 *draft-ietf-mppls-ldp-ipv6*, is present in the Hello message from the peer. This requirement is designed to address interoperability issues found with existing third-party LDP IPv4 implementations.


IP address and FEC distribution behavior is as follows:

- If the peer LSR sent the dual-stack capability TLV in the Hello message, then local IPv6 addresses are sent to the peer. The user can configure an address export policy to restrict which local IPv6 interface addresses are sent to the peer.
 - If the peer explicitly stated enabling of LDP IPv6 FEC type by including the IPv6 SAC TLV in the initialization message with the D-bit set to 0, then IPv6 FECs are also sent to the peer.
 - If the peer sent the dual-stack capability TLV in the Hello message, but explicitly stated disabling of LDP IPv6 FEC type by including the IPv6 SAC TLV in the initialization message with the D-bit set to 1, then IPv6 local addresses instead of IPv6 FECs are sent to the peer. The user can configure an address export policy to further restrict which local IPv6 interface addresses to send to the peer.
- If the peer did not send the dual-stack capability TLV in the Hello message, then no IPv6 addresses or IPv6 FECs are sent to that peer, regardless of the presence or not of the IPv6 SAC TLV in the initialization message. This case is added to prevent interoperability issues with some third-party LDP IPv4 implementations. The user can override the distribution defined by the initial Hello message by

explicitly configuring an address export policy and a FEC export policy to select IPv6 addresses and FECs to send to the peer.


The above behavior applies to LDP IPv4 and IPv6 addresses and FECs. The procedure is summarized in the flowchart diagrams in the following figures.

Figure 54: LDP IPv6 address and FEC distribution procedure



27907

Figure 55: LDP IPv6 address and FEC distribution procedure



27906

6.4.7 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 the LDP IPv4 or LDP IPv6 sessions comes up, IGP starts the sync timer. When the sync timer expires, the interface metric is restored to its configured value and all static routes with the **ldp-sync** option enabled are activated.

Given the above behavior, it is recommended that the user configure the sync timer to a value that allows enough time for both the LDP IPv4 and LDP IPv6 sessions to come up.

6.4.8 BFD operation

The operation of BFD over an 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. Tracking is required because LDP can 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. If any of the BFD tracking sessions fail, the LFA backup NHLFE for the FEC is activated, or, if there is no FRR backup, the FEC is re-resolved.

The user can configure tracking with only an IPv4 BFD session, only an IPv6 BFD session, or with both using the **config>router>ldp>if-params>if>bfd-enable [ipv4] [ipv6]** command.

This command provides flexibility in case the user does not need to track both Hello adjacency and the next hops of FECs.

For example, if the user configures **bfd-enable ipv6** only to save on the number of BFD sessions, then 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 and the IPv6 Hello adjacency is not affected 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, then all IPv4 and IPv6 prefix FECs are updated.

6.4.9 Services using SDP with an LDP IPv6 FEC

The 7705 SAR supports SDPs of type LDP with far-end options using IPv6 addresses. 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 admin groups are not supported with an SDP using an LDP IPv6 FEC, and the attempt to assign them is blocked in CLI.

Services that use the LDP control plane (such as T-LDP VPLS and R-VPLS, VLL, and IES/VRN 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. By default, the spoke SDP for these services binds to an SDP that uses an LDP IPv6 FEC that matches the prefix of the far end address.

In addition, the IPv6 PW control word is supported with data plane packets and VCCV OAM packets. Hash label is also supported with the above services, including the signaling and negotiation of hash label support using T-LDP (Flow sub-TLV) with the LDP IPv6 control plane. Finally, network domains are supported in VPLS.

6.4.10 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 VC-ID that is configured in the mirror destination service at the mirror source node. The **far-end** option is not supported with an IPv6 address.

6.4.10.1 Configuration at mirror source node

Use the following rules and syntax to configure a spoke SDP at the mirror source node:

- The *sdp-id* must match an SDP that uses an LDP IPv6 FEC.
- Configuring the *egress-vc-label* is optional.

CLI syntax:

```
no spoke-sdp sdp-id:vc-id
spoke-sdp sdp-id:vc-id [create] egress
vc-label egress-vc-label
```

6.4.10.2 Configuration at mirror destination node

Use the following rules and syntax to configure mirror service at the mirror destination node:

- The **far-end ip-address** command is not supported with LDP IPv6 transport tunnel. The user must reference a spoke SDP using an LDP IPv6 SDP coming from mirror source node.
- In the **spoke-sdp sdp-id:vc-id** command, the *vc-id* should match that of the **spoke-sdp** configured in the **mirror-destination** context at the mirror source node.
- Configuring the *ingress-vc-label* is optional; both Static and T-LDP are supported.

CLI syntax:

```
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
```

Mirroring is also supported with the PW redundancy feature when the endpoint spoke SDP, including the ICB, is using an LDP IPv6 tunnel.

6.4.11 OAM support with LDP IPv6

MPLS OAM tools LSP ping and LSP trace can operate with LDP IPv6 and support the following:

- use of IPv6 addresses in the echo request and echo reply messages, including in DMAP TLV, as per RFC 4379
- use of LDP IPv6 prefix target FEC stack TLV, as per RFC 4379
- use of IPv6 addresses in the DDMAP TLV and FEC stack change sub-TLV, as per 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, as per RFC 4379
- 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 supports both LDP IPv4 and IPv6 target FEC stack TLVs. Specifically:

- The IP family (IPv4/IPv6) of the UDP/IP echo request message always matches the family of the LDP target FEC stack TLV as entered by the user in the **prefix** option.
- The **src-ip-address** option is extended to accept 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 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 DMAP/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 the experimental value of 65503 in the Router Alert Option in the echo request packet IPv6 header, as per RFC 5350. After a value is allocated by IANA for MPLS OAM as part of draft-ietf-mpls-oam-ipv6-rao, it is updated.

VCCV ping and VCCV trace for a single-hop PW support IPv6 PW FEC 128, as per 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. When the value of the Channel Type field is set to 0x57, it indicates that the Associated Channel carries an IPv6 packet, as per RFC 4385.

6.4.12 Interoperability

6.4.12.1 Interoperability with implementations compliant with draft-ietf-mpls-ldp-ipv6

The 7705 SAR 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 is so that a routable system IPv6 address can be used by default to bring up the LDP task on the router and establish link LDP and T-LDP sessions to other LSRs. More importantly, using a 128-bit LSR ID allows for the establishment of control plane-independent LDP IPv4 and IPv6 sessions between two LSRs over the same interface or different set of interfaces because each session uses a unique LSR ID (32-bit for IPv4 and 128-bit for IPv6).

The 7705 SAR LDP implementation does not interoperate with a system using a 32-bit LSR ID (as defined in *draft-ietf-mpls-ldp-ipv6*) to establish an IPv6 LDP session. The latter specifies that an LSR can send both IPv4 and IPv6 Hello messages over an interface, allowing the system to 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 7705 SAR LDP implementation interoperates with systems 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 7705 SAR 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 thus for 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 7705 SAR implementation has a separate session for each IP family between two LSRs and, as such, this TLV is used to specify the family preference and to indicate that the system supports resolving IPv6 FECs over an IPv4 LDP session.

6.4.12.2 Interoperability with implementations compliant with RFC 5036 for IPv4 LDP control plane only

Some third-party LDP implementations are compliant with RFC 5036 for LDP IPv4 but are not compliant with RFC 5036 for handling IPv6 address or IPv6 FECs over an LDP IPv4 session.

An LSR based on the 7705 SAR in a LAN with a broadcast interface can peer with any third-party LSR, including those that are incapable of handling IPv6 address or IPv6 FECs over an LDP IPv4 session. When the 7705 SAR uses the IPv4 LDP control plane to advertise IPv6 addresses or IPv6 FECs to that peer, it can cause the IPv4 LDP session to go down.

To address this issue, *draft-ietf-mpls-ldp-ipv6* modifies RFC 5036 and requires compliant systems to check for the dual-stack capability TLV in the IPv4 Hello message from the peer. If the peer does not advertise this TLV, the LSR does not send IPv6 addresses and FECs to that peer. The 7705 SAR supports advertising and resolving IPv6 prefix FECs over an LDP IPv4 session using a 32-bit LSR ID in compliance with *draft-ietf-mpls-ldp-ipv6*.


6.4.13 Upgrading from IPv4 to IPv6

For a smooth transition from IPv4 to IPv6, the following steps are recommended:

1. Create a new IPv6 interface in the 7750 SR management VPRN.
2. Configure a new Layer 3 Spoke SDP configured with LDPv6 and a far-end IPv6 address, and assign it to the new IPv6 interface.
3. On the 7705 SAR, in the management Epipe create an endpoint object and assign the endpoint to the existing IPv4 PW. Ensure there is no traffic lost during this step.
4. On the 7705 SAR, create a new SDP with LDPv6 and a far-end IPv6 address.
5. On the 7705 SAR, within the management Epipe, assign the new SDP to a spoke SDP with the same endpoint as the IPv4 spoke SDP.
6. On the 7750 SAR, shut down the IPv4 interface.
7. On the 7705 SAR, start IPv6 traffic and ensure reachability to 7705 via the IPv6 SDP.
8. Remove the IPv4 SDP and spokes from the 7705 SAR Epipe and 7750 SR VPRN.

The following figure shows an example of a network ensuring a smooth upgrade from IPv4 to IPv6 with PW redundancy.


Figure 56: Smooth management transition from IPv4 to IPv6



6.5 LDP process overview

The following figure displays the process to provision basic LDP parameters.

Figure 57: LDP configuration and implementation



21820

6.6 Configuration notes

See the 7705 SAR Services Guide for information about signaling.

6.7 Configuring LDP with CLI

This section provides information to configure LDP using the command line interface.

Topics in this section include:

- [LDP configuration overview](#)
- [Basic LDP configuration](#)
- [Common configuration tasks](#)
- [LDP configuration management tasks](#)

6.8 LDP configuration overview

When the 7705 SAR 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*.

LDP must be enabled in order for signaling to be used to obtain the ingress and egress labels in frames transmitted and received on the service destination point (SDP). When signaling is off, labels must be manually configured when the SDP is bound to a service.

6.9 Basic LDP configuration

The LDP protocol instance is created in the **no shutdown** (enabled) state.

The following example displays the default LDP configuration output.

```
ALU-1>config>router>ldp# info
-----
      interface-parameters
      exit
      targeted-session
      exit
-----
ALU-1>config>router>ldp#
```

6.10 Common configuration tasks

This section provides a brief overview of the following tasks to configure LDP and provides the CLI commands:

- [Enabling LDP](#)
- [Configuring graceful restart helper parameters](#)
- [Applying import and export policies](#)
- [Configuring interface parameters](#)

- [Specifying targeted session parameters](#)
- [Specifying peer parameters](#)
- [Configuring LDP support for multicast VPN](#)
- [Configuring LDP support for LDP-to-SR stitching](#)
- [Enabling LDP signaling and services](#)

6.10.1 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 CLI syntax to enable LDP on a 7705 SAR router:

CLI syntax:

```
ldp
```

Example:

```
config>router# ldp
```

The following example displays the enabled LDP configuration output.

```
ALU-1>config>router# info
-----
...
#-----
echo "LDP Configuration"
#-----
      ldp
        interface-parameters
        exit
        targeted-session
        exit
      exit
-----
...
ALU-1>config>router#
```

6.10.2 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 7705 SAR recovery is self-contained and relies on information stored internally to self-heal.

Maximum recovery time is the time (in seconds) that 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) that the LSR is willing to retain its MPLS forwarding state. The time should be long enough to allow the neighboring LSRs to resynchronize all the LSPs in a graceful manner, without creating congestion in the LDP control plane.

Use the following syntax to configure graceful restart parameters:

CLI syntax:

```
config>router>ldp
[no] graceful-restart
[no] maximum-recovery-time interval
[no] neighbor-liveness-time interval
```

Example:

```
config>router>ldp
config>router>ldp# graceful-restart
config>router>ldp>graceful-restart# maximum-recovery-
time 120
config>router>ldp>graceful-restart# neighbor-liveness-
time 60
config>router>ldp# exit
config>router#
```

The following example displays the import policy configuration output.

```
ALU-1>config>router>ldp>graceful-restart# info
-----
maximum-recovery-time 120
neighbor-liveness-time 60
-----
ALU-1>config>router>ldp>graceful-restart#
```

6.10.3 Applying import and export policies

Inbound filtering (import policy) allows a route policy to control the label bindings that 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 the following:

- neighbor – match on bindings received from the specified peer
- prefix-list – match on bindings with the specified prefix or prefixes

Outbound filtering (export policy) allows a route policy to control the label bindings advertised by the LSR to its peers. Label bindings can be filtered based on the following:

- all – all local subnets by specifying **direct** as the match protocol
- prefix-list – match on bindings with the specified prefix/prefixes

Import or export policies must already exist before they are applied to LDP. Policies are configured in the **config>router>policy-options** context. See the "Route Policies" section in the 7705 SAR Router Configuration Guide for details.



Note:

- The 7705 SAR supports a specific number of labels, which varies by platform and software release. If the number of labels is exceeded for a specific protocol (for example, LDP or RSVP), a log message will appear by default in logs 99 and 100. The log message states the affected protocol and the label count that was exceeded. For example: "mpls_label_ilm_helper: XXXX XXX XXXX limit reached max obj count of YYYY".

- For the LDP protocol, when the label count is exceeded, LDP sessions will be shut down and all labels will be removed. To recover the LDP sessions, perform a **shutdown/no shutdown** combination of commands in the **config>router>ldp** context.

Use the following CLI syntax to apply import or export policies:

CLI syntax:

```
config>router>ldp
  import policy-name [policy-name...(up to 5 max)]
  export policy-name [policy-name...(up to 5 max)]
```

Example:

```
config>router>ldp
config>router>ldp# import LDP-import
config>router>ldp# export LDP-export
config>router>ldp# exit
config>router#
```

The following example displays the import and export policy configuration output.

```
ALU-1>config>router>ldp# info
-----
      export "LDP-export"
      import "LDP-import"
      interface-parameters
      exit
      targeted-session
      exit
-----
```

6.10.4 Configuring interface parameters

Use the following CLI syntax to configure LDP interface parameters:

CLI syntax:

```
config>router# ldp
  interface-parameters
    hello timeout factor
    interface ip-int-name
      hello timeout factor
      keepalive timeout factor
      local-lsr-id {system | interface}
      transport-address {system | interface}
      no shutdown
    keepalive timeout factor
    transport-address {system|interface}
```

Example:

```
config>router# ldp
config>router>ldp# interface-parameters
config>router>ldp>if-params# interface to-104
config>router>ldp>if-params>if# hello 15 3
config>router>ldp>if-params>if# local-lsr-id system
config>router>ldp>if-params>if# no shutdown
config>router>ldp>if-params>if# exit
config>router>ldp>if-params# exit
```

```
config>router>ldp#
```

The following example displays the LDP interface parameter configuration output.

```
ALU-1>config>router>ldp# info
-----
import "LDP-import"
interface-parameters
  hello 15 3
  keepalive 30 3
  interface "to-104"
    hello 15 3
    keepalive 30 3
    local-lsr-id system
    no shutdown
  exit
exit
targeted-session
exit
no shutdown
-----
ALU-1>config>router>ldp#
```

6.10.5 Specifying targeted session parameters

Use the following CLI syntax to specify targeted session parameters:

CLI syntax:

```
config>router# ldp
targeted-session
  disable-targeted-session
  hello timeout factor
  keepalive timeout factor
  peer ip-address
  bfd-enable
  hello timeout factor
  keepalive timeout factor
  local-lsr-id interface-name
  no shutdown
```

Example:

```
config>router# ldp
config>router>ldp# targeted-session
config>router>ldp>targ-session# bfd-enable
config>router>ldp>targ-session# hello 5000 255
config>router>ldp>targ-session# keepalive 5000 255
config>router>ldp>targ-session# peer 10.10.10.104
config>router>ldp>targ-session>peer# hello 2500 100
config>router>ldp>targ-session>peer# keepalive 15 3
config>router>ldp>targ-session>peer# local-lsr-id to-104
config>router>ldp>targ-session>peer# no shutdown
config>router>ldp>targ-session>peer# exit
config>router>ldp>targ-session# exit
config>router>ldp#
```

The following example displays the LDP targeted session configuration output.

```
ALU-1>config>router>ldp# info
```

```

-----
import "LDP-import"
interface-parameters
    hello 15 3
    keepalive 30 3
    interface "to-104"
        hello 15 3
        keepalive 30 3
        no shutdown
    exit
exit
targeted-session
    hello 5000 255
    keepalive 5000 255
    peer 10.10.10.104
        hello 2500 100
        keepalive 15 3
        local-lsr-id "to-104"
    exit
exit
-----

```

6.10.6 Specifying peer parameters

Use the following CLI syntax to specify LDP peer parameters:

CLI syntax:

```

config>router# ldp
peer-parameters
    peer ip-address
        auth-keychain name
        authentication-key {authentication-key | hash-key} [hash |
hash2]

```

Example:

```

config>router# ldp
config>router>ldp# peer-parameters
config>router>ldp>peer-params# peer 10.10.10.104
config>router>ldp>peer-params>peer$ authentication-key
testuser
config>router>ldp>peer-params>peer$ exit

```

The following example displays the LDP peer parameters configuration output.

```

ALU-1>config>router>ldp# info
-----
import "LDP-import"
graceful-restart
exit
import "LDP-import"
peer-parameters
    peer 10.10.10.104
        authentication-key "nGjXyHQtcGhxbBm.kDeYdzSmPZy9KK03" hash2
    exit
exit
interface-parameters
    interface "test"
    exit
    interface "to-104"

```

```

        hello 15 3
    exit
exit
targeted-session
    hello 5000 255
    keepalive 5000 255
    peer 10.10.10.104
        hello 2500 100
        keepalive 15 3
    exit
exit
-----
ALU-1>config>router>ldp#

```

6.10.7 Configuring LDP support for multicast VPN

For LDP support for multicast VPN (MVPN), configure the **multicast-traffic**, **mp-mbb-time**, and **mcast-upstream-frr** commands.

The following example displays the LDP MVPN configuration output.

```

*A:SaRAx Dut-D>config>router>ldp# info detail
-----
    no aggregate-prefix-match
    no export
    no fast-reroute
    no import
    no graceful-restart
    mcast-upstream-frr
    mp-mbb-time 5
    no tunnel-down-damp-time
    interface-parameters
        hello 15 3
        keepalive 30 3
        transport-address system
        interface "mcast_if"
            no bfd-enable
            no hello
            no keepalive
            no local-lsr-id
            multicast-traffic enable
            no transport-address
            no shutdown
        exit
    exit
    targeted-session
        no disable-targeted-session
        hello 45 3
        keepalive 40 4
    exit
    no shutdown
-----
*A:SaRAx Dut-D>config>router>ldp#

```

6.10.8 Configuring LDP support for LDP-to-SR stitching

Configure the **export-tunnel-table** command using the following CLI syntax to support LDP-to-SR stitching.

CLI syntax:

```
config>router# ldp
export-tunnel-table policy-name [policy-name...(up to 5 max)]
```



Note: The specified policy name must be the same as the policy-statement name defined when configuring the route policy options for LDP-to-SR stitching. See the 7705 SAR Router Configuration Guide, "Configuring LDP-to-Segment Routing Stitching Policies".

Example:

```
config>router>ldp# export-tunnel-table "export-SR"
config>router>ldp# exit
```

The following example displays the LDP-to-SR stitching LDP configuration output.

```
*A:NOK-1 Dut-B>config>router>ldp# info
-----
export-tunnel-table "export-SR"
...
exit
...
-----
*A:NOK-1 Dut-B>config>router>ldp#
```

6.10.9 Enabling 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 destination points (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 triggers session establishment. The SDP's signaling default enables **tl**dp. The SDP uses the targeted-session parameters configured in the **config>router>ldp> targeted-session** context.

The **service>sdp>ldp** and **router>lsp** commands are mutually exclusive; you can either specify an LSP or enable an LDP. There cannot be two methods of transport in a single SDP.

To enable LDP on the SDP when an LSP is already specified, the LSP must be removed from the configuration using the **no lsp lsp-name** command. For more information about configuring SDPs, see the 7705 SAR Services Guide.

Use the following CLI syntax to enable LDP on an MPLS SDP:

CLI syntax:

```
config>service>sdp#
ldp
signaling {off|tl}dp}
```

The following example displays an SDP configuration output with the signaling default **tl dp** enabled.

```
ALU-1>config>service>sdp# info detail
-----
      description "MPLS: to-99"
      far-end 10.10.10.99
      ldp
      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
-----
ALU-1>config>service>sdp#
```

6.11 LDP configuration management tasks

This section discusses the following LDP configuration management tasks:

- [Disabling LDP](#)
- [Modifying targeted session parameters](#)
- [Modifying interface parameters](#)

6.11.1 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 CLI syntax to disable LDP:

CLI syntax:

```
no ldp
shutdown
```

Example:

```
config>router# ldp
config>router>ldp# shutdown
config>router>ldp# exit
config>router# no ldp
```

6.11.2 Modifying targeted session parameters

You can modify targeted session parameters without shutting down entities. However, for any LDP timers (hello or keepalive timers), the changes do not take effect until a **shutdown** or **no shutdown** command is performed on the LDP session.

The **no** form of a **targeted-session** parameter command reverts modified values back to the default.

The following example displays the CLI syntax to revert targeted session parameters back to the default values.

Example:

```
config>router# ldp
config>router>ldp# targeted-session
config>router>ldp>targeted# no disable-targeted-session
config>router>ldp>targeted# no hello
config>router>ldp>targeted# no keepalive
config>router>ldp>targeted# shutdown
config>router>ldp>targeted# no shutdown
config>router>ldp>targeted# no peer 10.10.10.99
```

The following example displays the default value output.

```
ALU-1>config>router>ldp>targeted# info detail
-----
no disable-targeted-session
hello 45 3
keepalive 40 4
-----
ALU-1>config>router>ldp>targeted#
```

6.11.3 Modifying interface parameters

You can modify LDP interface parameters without shutting down entities. However, at the global timer configuration level (**ldp>interface-parameters**), the **hello** and **keepalive** parameter modifications do not take effect until a **shutdown** or **no shutdown** command is performed on the LDP session. At the interface timer configuration level (**ldp>interface-parameters>interface**), any changes to the **keepalive** parameter do not take effect until a **shutdown/no shutdown** command is performed on the LDP session. For all other parameters, the changes take effect immediately.

Individual parameters cannot be deleted. The **no** form of an **interface-parameter** command changes modified values back to the defaults.

The following example displays the CLI syntax to change interface parameters back to the default values.

Example:

```
config>router# ldp
config>router>ldp>interface-parameters
config>router>ldp>if-params# no hello
config>router>ldp>if-params# interface to-104
config>router>ldp>if-params>if# no keepalive
config>router>ldp>if-params>if# no transport-address
config>router>ldp>if-params>if# shutdown
config>router>ldp>if-params>if# no shutdown
config>router>ldp>if-params>if# exit
config>router>ldp>if-params# exit
config>router>ldp# shutdown
config>router>ldp# no shutdown
```

The following example displays the default value output.

```
ALU-1>config>router>ldp>if-params# info detail
-----
hello 15 3
```



```
        keepalive 30 3
        no transport-address
-----
ALU-1>config>router>ldp>params#
```

6.12 LDP command reference

6.12.1 Command hierarchies

- [LDP commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Monitor commands](#)
- [Debug commands](#)

6.12.1.1 LDP commands

```

config
- router router-name
- [no] ldp
- [no] aggregate-prefix-match
- prefix-exclude policy-name [policy-name...(up to 5 max)]
- no prefix-exclude
- [no] shutdown
- egress-statistics
- [no] fec-prefix ip-prefix[/mask]
- accounting-policy policy-id
- no accounting-policy
- [no] collect-stats
- [no] shutdown
- export policy-name [policy-name...(up to 5 max)]
- no export
- export-tunnel-table policy-name [policy-name...(up to 5 max)]
- no export-tunnel-table
- fast-reroute [backup-sr-tunnel]
- no fast-reroute
- fec-originate ip-address/mask [advertised-label in-label] [swap-label out-label]
interface interface-name
- fec-originate ip-address/mask [advertised-label in-label] next-hop ip-address
[swap-label out-label]
- fec-originate ip-address/mask [advertised-label in-label] next-hop ip-address
[swap-label out-label] interface interface-name
- fec-originate ip-address/mask [advertised-label in-label] pop
- no fec-originate ip-address/mask interface interface-name
- no fec-originate ip-address/mask next-hop ip-address
- no fec-originate ip-address/mask next-hop ip-address interface interface-name
- no fec-originate ip-address/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)]
- no import
- interface-parameters
- interface ip-int-name [dual-stack]
- [no] interface ip-int-name
- bfd-enable [ipv4] [ipv6]

```

```

- [no] bfd-enable
- [no] ipv4
  - fec-type-capability
    - p2mp-ipv4 {enable | disable}
    - p2mp-ipv6 {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}
  - no local-lsr-id
  - [no] shutdown
  - transport-address {system | interface}
  - no transport-address
- [no] ipv6
  - fec-type-capability
    - p2mp-ipv4 {enable | disable}
    - p2mp-ipv6 {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}
  - no local-lsr-id
  - [no] shutdown
  - transport-address {system | interface}
  - no transport-address
- ipv4
  - hello timeout factor
  - no hello
  - keepalive timeout factor
  - no keepalive
  - transport-address {system | interface}
  - no transport-address
- ipv6
  - hello timeout factor
  - no hello
  - keepalive timeout factor
  - no keepalive
  - transport-address {system | interface}
  - no transport-address
- [no] legacy-ipv4-lsr-interop
- [no] mcast-upstream-frr
- mp-mbb-time interval
- no mp-mbb-time
- session-parameters
  - [no] peer ip-address
  - 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-type-capability
    - p2mp {enable | disable}
    - prefix-ipv4 {enable | disable}
    - prefix-ipv6 {enable | disable}
- [no] shutdown
- targeted-session
  - [no] disable-targeted-session
  - ipv4
    - hello timeout factor

```

```

- no hello
- keepalive timeout factor
- no keepalive
- ipv6
- hello timeout factor
- no hello
- keepalive timeout factor
- no keepalive
- [no] peer ip-address
- [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
- [no] peer-transport ip-address
- auth-keychain name
- no auth-keychain
- authentication-key {authentication-key | hash-key} [hash | hash2]
- no authentication-key
- no authentication-key
- tunnel-down-damp-time seconds
- no tunnel-down-damp-time

```

6.12.1.2 Show commands

The show commands in this section are organized into the following areas:

- [Show LDP commands](#)
- [Show LDP bindings commands](#)

6.12.1.2.1 Show LDP commands

```

show
- router [router-instance]
- ldp
- discovery [state state] [detail | summary] [adjacency-type type] [session ip-
addr[label-space]]
- discovery [state state] [detail | summary] [adjacency-type type] [family]
- discovery interface [ip-int-name] [state state] [detail | summary] [session ip-
addr[label-space]] [family]
- discovery peer [ip-address] [state state] [detail | summary] [session ip-
addr[label-space]]
- fec-egress-stats [ip-prefix/mask]
- fec-egress-stats [active] [family]
- fec-originate [ip-address/mask] [operation-type]
- fec-originate [operation-type] [family]
- interface [ip-int-name | ip-address] [detail] [family]
- parameters
- session [ip-addr[:label-space]] local-addresses [sent | rcv] ip-addr ip-address
- session [ip-addr[:label-space]] [session-type] [state state] [detail | summary]
- session [ip-addr[:label-space]] local-addresses [sent | rcv] [family]
- session [ip-addr[:label-space]] statistics [packet-type] [session-type]

```

```

- session statistics [packet-type] [session-type] [family]
- session [session-type] [state state] [detail | summary] [family]
- session-parameters [family]
- session-parameters [peer-ip-address]
- status
- targ-peer [ip-address] [detail]
- targ-peer [detail] [family]
- tcp-session-parameters [family]
- tcp-session-parameters [keychain keychain]
- tcp-session-parameters transport-peer-ip-address

```

6.12.1.2.2 Show LDP bindings commands

```

show
- router
- ldp
- bindings
- active detail [family] [egress-if port-id]
- active detail [family] [egress-lsp tunnel-id]
- active detail [egress-nh ip-address] [family]
- active egress-if port-id [summary | detail] [family]
- active egress-lsp tunnel-id [summary | detail] [family]
- active egress-nh [family] [summary | detail] ip-address
- active ipv4 [summary | detail] [egress-if port-id]
- active ipv4 [summary | detail] [egress-lsp tunnel-id]
- active ipv4 [summary | detail] [egress-nh ip-address]
- active ipv6 [summary | detail] [egress-if port-id]
- active ipv6 [summary | detail] [egress-nh ip-address]
- active ipv6 [summary | detail] [egress-lsp tunnel-id]
- active p2mp p2mp-id identifier root ip-address [summary | detail] [egress-
if port-id]
- active p2mp p2mp-id identifier root ip-address [summary | detail] [egress-
lsp tunnel-id]
- active p2mp p2mp-id identifier root ip-address [summary | detail] [egress-
nh ip-address]
- active p2mp [family] [summary | detail] [egress-if port-id] [opaque-
type opaque-type]
- active p2mp [family] [summary | detail] [egress-lsp tunnel-id] [opaque-
type opaque-type]
- active p2mp [family] [summary | detail] [egress-nh ip-address] [opaque-
type opaque-type]
- active p2mp source ip-address group mcast-address root ip-address [summary |
detail] [egress-if port-id] inner-root ip-address
- active p2mp source ip-address group mcast-address root ip-address [summary |
detail] [egress-lsp tunnel-id] inner-root ip-address
- active p2mp source ip-address group mcast-address root ip-address [summary |
detail] [egress-nh ip-address] inner-root ip-address
- active p2mp source ip-address group mcast-address root ip-address [rd rd]
[summary | detail] [egress-if port-id]
- active p2mp source ip-address group mcast-address root ip-address [rd rd]
[summary | detail] [egress-lsp tunnel-id]
- active p2mp source ip-address group mcast-address root ip-address [rd rd]
[summary | detail] [egress-nh ip-address]
- active p2mp source ip-address group mcast-address [family] [summary | detail]
[innermost-root ip-address]
- active prefixes [family] [summary | detail] [egress-if port-id]
- active prefixes [family] [summary | detail] [egress-lsp tunnel-id]
- active prefixes [egress-nh ip-address] [family] [summary | detail]
- active prefixes prefix ip-prefix/ip-prefix-length [summary | detail] [egress-
if port-id]

```

```

- active prefixes prefix ip-prefix/ip-prefix-length [summary | detail] [egress-
lsp tunnel-id]
- active prefixes prefix ip-prefix/ip-prefix-length [egress-nh ip-address]
[summary | detail]
- active summary [family] [egress-if port-id]
- active summary [family] [egress-lsp tunnel-id]
- active summary [egress-nh ip-address] [family]

show
- router
- ldp
- bindings
- detail [session ip-addr [label-space]] [family]
- ipv4 [session ip-addr[label-space]] [summary | detail]
- ipv6 [session ip-addr[label-space]] [summary | detail]
- label-type start-label start-label [end-label end-label] label-type [family]
- p2mp p2mp-id identifier root ip-address [session ip-addr [label-space]]
[summary | detail]
- p2mp [session ip-addr [label-space]] [family] [summary | detail] [opaque-
type opaque-type]
- p2mp source ip-address group mcast-address root ip-address [session ip-
addr [label-space]] [family] [summary | detail] inner-root ip-address
- p2mp source ip-address group mcast-address root ip-address [rd rd]
[session ip-addr [label-space]] [summary | detail]
- p2mp source ip-address group mcast-address [session ip-addr [label-space]]
[family] [summary | detail] [innermost-root ip-address]
- prefixes prefix ip-prefix/ip-prefix-length [summary | detail] [session ip-
addr[label-space]]
- prefixes [family] [summary | detail] [session ip-addr[label-space]]
- services vc-type vc-type saii global-id:prefix:ac-id [256 chars max]taii
agi agi [detail] [service-id service-id] [session ip-addr[label-space]]
- services vc-type vc-type agi agi [detail] [service-id service-id]
[session ip-addr[label-space]]
- services [vc-type vc-type] [svc-fec-type] [detail] [service-id service-id]
[session ip-addr[label-space]]
- services vc-type vc-type vc-id vc-id [detail] [service-id service-id]
[session ip-addr[label-space]]
- session [family] [summary | detail] ip-addr[label-space]
- summary [session ip-addr[label-space]] [ipv4 | ipv6]

```

6.12.1.3 Monitor commands

```

monitor
- router
- ldp
- fec-egress-stats ip-prefix/mask [interval seconds] [repeat repeat] [absolute |
rate]
- session ldp-id[ldp-id...(up to 5 max)] [interval seconds] [repeat repeat]
[absolute | rate]
- statistics [interval seconds] [repeat repeat] [absolute | rate]

```

6.12.1.4 Clear commands

```

clear
- router [router-instance]
- ldp
- fec-egress-statistics [ip-prefix/mask]
- instance
- interface ip-int-name [statistics]

```

```
- peer ip-address [statistics]
- session [ip-addr[:label-space]] [statistics]
- statistics
```

6.12.1.5 Debug commands

```
[no] debug
- router [router-instance]
- [no] ldp
- [no] interface interface-name
- [no] event
- [no] messages
- [no] packet
- hello [detail]
- no hello
- [no] 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
```

6.12.2 Command descriptions

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

6.12.2.1 Configuration commands

- [Generic commands](#)
- [LDP global commands](#)
- [Interface parameters commands](#)
- [Session parameters commands](#)
- [Targeted session commands](#)

6.12.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

```
config>router>ldp
config>router>ldp>egress-statistics
config>router>ldp>if-params>interface
config>router>ldp>if-params>if>ipv4
config>router>ldp>if-params>if>ipv6
config>router>ldp>targ-session>peer
config>router>ldp>aggregate-prefix-match
```

Description

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

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

The **no** form of this command administratively enables an entity.

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

Default

no shutdown

6.12.2.1.2 LDP global commands

ldp

Syntax

[no] ldp

Context

config>router

Description

This command enables the context to configure an LDP protocol instance.

When an LDP instance is created, the protocol is enabled (in the **no shutdown** state). To suspend the LDP protocol, use the **shutdown** command. Configuration parameters are not affected.

The no form of the command deletes the LDP protocol instance, removing all associated configuration parameters. The LDP instance must first be disabled with the **shutdown** command before being deleted.

Default

n/a – LDP must be explicitly enabled

aggregate-prefix-match

Syntax

[no] aggregate-prefix-match

Context

config>router>ldp

Description

This command enables LDP to use the aggregate prefix match function instead of requiring an exact prefix match.

When this command is enabled and an LSR receives a FEC-label binding from an LDP neighbor for a prefix-address FEC element, FEC1, it installs the binding in the LDP FIB if:

- the routing table (RIB) contains an entry that matches FEC1. Matching can either be a longest IP match of the FEC prefix or an exact match.

- the advertising LDP neighbor is the next hop to reach FEC1

When the FEC-label binding has been installed in the LDP FIB, LDP programs an NHLFE entry in the egress data path 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 RIB, 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, this may mean that the LSR used as the next hop changes; if so, the NHLFE entry for that FEC must be changed.

When a prefix is removed from the RIB, 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, LDP must use it. This may mean that the LSR used as the next hop also changes; if so, the NHLFE entry for that FEC must be changed. If another match does not exist, the LSR removes the FEC binding and sends a label withdraw message to its LDP neighbors.

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 entry for the FEC elements.

The **no** form of this command disables the use of the aggregate prefix match function. LDP then only performs an exact prefix match 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

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 the prefix occurs in the routing table.

The **no** form of this command removes all policies from the configuration; therefore, no prefixes are excluded.

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 (such as #, \$, spaces), the entire string must be enclosed within double quotes.

egress-statistics

Syntax

egress-statistics

Context

config>router>ldp

Description

This command enters the context to enable egress data path statistics at the ingress LER for this FEC.

Default

n/a

fec-prefix

Syntax

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

Context

config>router>ldp>egress-statistics

Description

This command configures statistics in the egress data path at the ingress LER or LSR for an LDP FEC. The user must also execute the **no shutdown** command in this context to enable statistics collection.

The **no** form of this command disables the statistics in the egress data path and removes the accounting policy association from the LDP FEC.

Default

n/a

Parameters

ip-prefix[/mask]

the IP prefix and prefix length associated with the prefix FEC

accounting-policy

Syntax

accounting-policy *policy-id*
no accounting-policy

Context

config>router>ldp>egr-stats>fec-prefix

Description

This command associates an accounting policy with an LDP FEC. Only one accounting policy at a time can be associated with an LDP FEC on a particular node.

An accounting policy must first be configured in the **config>log>accounting-policy** context before it can be associated; otherwise an error message is generated.

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

Default

no accounting policy

Parameters

policy-id
the accounting policy ID
Values 1 to 99

collect-stats

Syntax

[no] collect-stats

Context

config>router>ldp>egr-stats>fec-prefix

Description

This command enables accounting and statistical data collection.

The collected statistic counters can be retrieved via **show** and **monitor** commands or with the SNMPv3 interface. The counters can be saved to an accounting file if the specific statistics collection record is included in the default accounting policy or in a user-defined accounting policy.

When the **no** form of this command is issued, statistics are still accumulated by the forwarding engine; however, the CPU does not obtain the results and write them to the accounting file. If a subsequent **collect-stats** command is issued, then the counters written to the accounting file also includes all the traffic that went through while the **no collect-stats** command was in effect.

Default

no collect-stats

export

Syntax

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

no export

Context

config>router>ldp

Description

This command specifies export route policies that determine which routes are exported to LDP neighbors. Configuring an export policy allows the LSR (Label Switch Router) to advertise addresses other than the system IP address. Policies are configured in the **config>router>policy-options** context. See the "Route Policies" section in the 7705 SAR Router Configuration Guide.

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

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

The **no** form of the 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 long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, spaces), the entire string must be enclosed within double quotes.

export-tunnel-table

Syntax

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

no export-tunnel-table

Context

config>router>ldp

Description

This command enables exporting of SR tunnels from the TTM into LDP (IGP) for the purpose of stitching an LDP FEC to an SR tunnel for the same destination IPv4 /32 IS-IS prefix.

The **no** form of the command disables the exporting of SR tunnels to LDP.

Default

no export-tunnel-table

Parameters

policy-name

the export-tunnel-table route policy name; must be an existing policy-statement name

fast-reroute

Syntax

fast-reroute [**backup-sr-tunnel**]

no fast-reroute

Context

config>router>ldp

Description

This command enables LDP fast reroute (FRR). LDP FRR provides local protection for an LDP FEC by precalculating and downloading a primary and a backup NHLFE for the FEC to the LDP FIB.

This command is limited to IPv4 /32 prefixes in both LDP and SR.

When LDP FRR is enabled and an LFA backup next hop exists for the FEC prefix in the RTM, or for the longest prefix the FEC prefix matches to when the [aggregate-prefix-match](#) command is enabled, LDP programs the data path with both a primary NHLFE and a backup NHLFE for each next hop of the 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
- an LDP session to a peer goes down because the Hello timer or keepalive timer has expired over an interface
- the TCP connection used by a link LDP session to a peer goes down

The [tunnel-down-damp-time](#) command, 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 can switch to the LFA next hop while the IGP (OSPF or IS-IS) is still using the primary next hop. As well, when the interface for the previous primary next hop is restored, the 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 with the **config>router>if>ldp-sync-timer** command (see the 7705 SAR Router Configuration Guide, "IP Router Command Reference", for information about configuring the **ldp-sync-timer**).

The **backup-sr-tunnel** option allows an SR tunnel to be used as a remote LFA or TI-LFA backup tunnel next hop by an LDP FEC. Before this option can be used, the LDP-to-SR stitching feature must be enabled. See [LDP-to-segment routing stitching for IPv4 /32 prefixes \(IS-IS\)](#) for more information about this feature.

The **no** form of this command disables LDP FRR.

Default

no fast-reroute

Parameters

backup-sr-tunnel

allows an SR tunnel to be used as a remote LFA or TI-LFA backup tunnel next hop by an LDP FEC

fec-originate

Syntax

fec-originate *ip-address/mask* [**advertised-label** *in-label*] [**swap-label** *out-label*] **interface** *interface-name*

fec-originate *ip-address/mask* [**advertised-label** *in-label*] **next-hop** *ip-address* [**swap-label** *out-label*]

fec-originate *ip-address/mask* [**advertised-label** *in-label*] **next-hop** *ip-address* [**swap-label** *out-label*] **interface** *interface-name*

fec-originate *ip-address/mask* [**advertised-label** *in-label*] **pop**

no fec-originate *ip-address/mask* **interface** *interface-name*

no fec-originate *ip-address/mask* **next-hop** *ip-address*

no fec-originate *ip-address/mask* **next-hop** *ip-address* **interface** *interface-name*

no fec-originate *ip-address/mask* **pop**

Context

config>router>ldp

Description

This command adds a FEC to the LDP prefix database with a specific label operation on the node.

Permitted operations are **swap** to originate a FEC for which the LSR is not egress or **pop** to originate a FEC for which the LSR is egress.

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, 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 the user-configured next hop for the swap operation must match one of the next hops in the route-table entry.

The **next-hop**, **advertised-label**, and **swap-label** parameters are optional. If a **next-hop** is configured but no **swap-label** is specified, the swap occurs with label 3 (implicit null), then the label is popped and the

packet is forwarded to the next hop. If the **next-hop** and **swap-label** parameters are configured, a regular swap occurs. If no parameters are specified, a pop and forwarding is performed.

Default

no fec-originate

Parameters

ip-address/mask

specifies the IP prefix and mask length

advertised label

specifies the label advertised to the upstream peer. If not configured, the label that is advertised should be from the label pool. If the configured static label is not available, the IP prefix is not advertised.

in-label

the LSR to swap the label. If configured, the LSR should swap the label with the configured *out-label*. If not configured, the default action is pop if the next-hop parameter is not defined.

Values 32 to 2047

out-label

the number of labels to send to the peer associated with this FEC

Values 16 to 1048575

interface-name

specifies the name of the interface that the label for the originated FEC is swapped to. For an unnumbered interface, this parameter is mandatory because there is no address for the next hop. For a numbered interface, it is optional.

next-hop *ip-address*

specifies the IP address of the next hop

pop

specifies to pop the label and transmit the packet

graceful-restart

Syntax

[no] graceful-restart

Context

config>router>ldp

Description

This command enables graceful restart helper.

The **no** form of the command disables graceful restart.

Default

graceful-restart

maximum-recovery-time

Syntax

maximum-recovery-time *interval*

no maximum-recovery-time

Context

config>router>ldp>graceful-restart

Description

This command configures the local maximum recovery time, which is the time (in seconds) that the sender of the TLV would like the receiver to wait, after detecting the failure of LDP communication with the sender.

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

Default

120

Parameters

interval

specifies the maximum length of recovery time, in seconds

Values 15 to 1800

neighbor-liveness-time

Syntax

neighbor-liveness-time *interval*

no neighbor-liveness-time

Context

config>router>ldp>graceful-restart

Description

This command configures the neighbor liveness time, which is the time (in seconds) that the LSR retains its MPLS forwarding state. The time should be long enough to allow the neighboring LSRs to resynchronize all the LSPs in a graceful manner, without creating congestion in the LDP control plane.

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

Default

120

Parameters*interval*

specifies the length of time, in seconds

Values 5 to 300**implicit-null-label****Syntax****[no] implicit-null-label****Context**

config>router>ldp

Description

This command enables the implicit null label option for all LDP FECs for which the router is the eLER.

The implicit null label is signaled by the eLER to the previous-hop LSR during FEC signaling by the LDP control protocol. When the implicit null label is signaled to the LSR, it pops the outer label before sending the MPLS packet to the eLER; this is known as penultimate hop popping.

The **no** form of the command disables the signaling of the implicit null label.

Default

no implicit-null-label

import**Syntax****import** *policy-name*[*policy-name* ...(up to 5 max)]**no import****Context**

config>router>ldp

Description

This command specifies import route policies that determine which routes are accepted from LDP neighbors. Policies are configured in the **config>router>policy-options** context. See the "Route Policies" section in the 7705 SAR Router Configuration Guide.

If no import policy is specified, LDP accepts all routes 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 specified names must already be defined.

The **no** form of the 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 long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, spaces), the entire string must be enclosed within double quotes.

hello

Syntax

hello *timeout factor*

no hello

Context

config>router>ldp>if-params>ipv4

config>router>ldp>if-params>ipv6

config>router>ldp>if-params>if>ipv4

config>router>ldp>if-params>if>ipv6

config>router>ldp>targ-session>ipv4

config>router>ldp>targ-session>ipv6

config>router>ldp>targ-session>peer

Description

This command configures the hold time. This is the time interval to wait before declaring a neighbor down. The *factor* parameter derives the hello interval.

Hold time is local to the system and is sent in the hello messages to the neighbor. Hold time cannot be less than three times the hello interval. The hold time can be configured globally (applies to all LDP interfaces) or per interface. The most specific value is used.

After an LDP session is being set up, the hold 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 **no** form of the command:

- at the interface-parameters and targeted-session levels, sets the **hello timeout** and the **hello factor** to the default values

- at the interface level, sets the **hello timeout** and the **hello factor** to the value defined under the interface-parameters level
- at the peer level, sets the **hello timeout** and the **hello factor** to the value defined under the targeted-session level

Default

The default value is dependent upon the CLI context. The following table lists the **hello timeout factor** default values.

Table 45: Hello timeout factor default values

Context	Timeout	Factor
config>router>ldp>if-params>ipv4 config>router>ldp>if-params>ipv6	15	3
config>router>ldp>if-params>interface>ipv4 config>router>ldp>if-params>interface>ipv6	Inherits values from interface-parameters context	
config>router>ldp>targ-session>ipv4 config>router>ldp>targ-session>ipv6	45	3
config>router>ldp>targ-session>peer	Inherits values from targeted-session IPv4 or IPv6 context	

Parameters

timeout

configures the time interval, in seconds, that LDP waits before declaring a neighbor down

Values 1 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

keepalive

Syntax

keepalive *timeout factor*
no keepalive

Context

config>router>ldp>if-params>ipv4
config>router>ldp>if-params>ipv6

```

config>router>ldp>if-params>if>ipv4
config>router>ldp>if-params>if>ipv6
config>router>ldp>targ-session>ipv4
config>router>ldp>targ-session>ipv6
config>router>ldp>targ-session>peer

```

Description

This command configures the time interval, in seconds, that LDP waits before tearing down the session. The **factor** parameter derives the keepalive interval.

If no LDP messages are exchanged for the configured time interval, the LDP session is torn down. Keepalive timeout is usually three times the keepalive interval. To maintain the session permanently, regardless of the activity, set the value to zero.

After an LDP session is set up, the keepalive timeout is negotiated to the lower of the two peers. After an operational value is agreed upon, the **keepalive factor** is used to derive the value of the keepalive interval.

The **no** form of the command:

- at the IPv4, IPv6, and targeted-session levels, sets the **keepalive timeout** and the **keepalive factor** to the default value
- at the IPv4 or IPv6 interface level, sets the **keepalive timeout** and the **keepalive factor** to the value defined under the corresponding **interface-parameters** level
- at the peer level, sets the **keepalive timeout** and the **keepalive factor** to the value defined under the **targeted-session** level

Default

The default value is dependent upon the CLI context. The following table lists the **keepalive timeout** factor default values.

Table 46: Keepalive timeout factor default values

Context	Timeout	Factor
config>router>ldp>if-params>ipv4 config>router>ldp>if-params>ipv6	30	3
config>router>ldp>if-params>interface>ipv4 config>router>ldp>if-params>interface>ipv6	Inherits values from interface-parameters context	
config>router>ldp>targ-session>ipv4 config>router>ldp>targ-session>ipv6	40	4
config>router>ldp>targ-session>peer	Inherits values from targeted-session IPv4 or IPv6 context	

Parameters*timeout*

configures the time interval, expressed in seconds, that LDP waits before tearing down the session

Values 1 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

legacy-ipv4-lsr-interop**Syntax**

[no] legacy-ipv4-lsr-interop

Context

config>router>ldp

Description

This command allows interoperability with third-party legacy IPv4 LSR implementations that do not comply with RFC 5036 with respect to the processing of Hello TLVs with the U-bit set.

The command is a global LDP configuration that disables the Nokia proprietary Interface Info TLV (0x3E05) in the Hello message sent to the peer. Disabling this Hello TLV also results in the non-generation of the Nokia proprietary Hello Adjacency Status TLV (0x3E06) because the Interface Info TLV is not sent.

In addition, this command disables the RFC 7552 standard dual-stack capability TLV (0x701) and the Nokia proprietary Adjacency capability TLV (0x3E07).

mcast-upstream-frr**Syntax**

[no] mcast-upstream-frr

Context

config>router>ldp

Description

This command enables the mLDP fast upstream switchover feature.

When this command is enabled and LDP is resolving an mLDP FEC received from a downstream LSR, it checks whether an ECMP next hop or an LFA next hop to the root LSR node exists. If LDP finds one, it programs a primary ILM on the interface corresponding to the primary next hop and a backup ILM on the interface corresponding to the ECMP or LFA next hop. Then, LDP sends the corresponding labels to

both upstream LSR nodes. Under normal operation, the primary ILM accepts packets while the backup ILM drops them. If the interface or the upstream LSR of the primary ILM goes down, causing the LDP session to go down, the backup ILM starts accepting packets.

To make use of the ECMP next hop, the user must configure the **ecmp** value in the system to at least "2", using the following command:

```
config>router>ecmp
```

To make use of the LFA next hop, the user must enable LFA using the following commands (as needed):

```
config>router>isis>loopfree-alternates
```

```
config>router>ospf>loopfree-alternates
```

Enabling the IP FRR or LDP FRR feature is not strictly required because LDP only needs to know the location of the alternate next hop to the root LSR so it can send the Label Mapping message to program the backup ILM at the initial signaling of the tree. Therefore, enabling the LFA option is sufficient. However, if unicast IP and LDP prefixes need to be protected, then these features and the mLDP fast upstream switchover can be enabled concurrently.

The mLDP FRR fast switchover relies on the fast detection of a loss of an LDP session to the upstream peer to which the primary ILM label had been advertised. It is strongly recommended that the following be performed:

1. Enable BFD on all LDP interfaces to upstream LSR nodes. When BFD detects the loss of the last adjacency to the upstream LSR, it immediately brings down the LDP session, which causes the CSM to activate the backup ILM.
2. If there is a concurrent T-LDP adjacency to the same upstream LSR node, enable BFD on the T-LDP peer in addition to enabling it on the interface.
3. Enable the **ldp-sync-timer** option on all interfaces to the upstream LSR nodes. If an LDP session to the upstream LSR to which the primary ILM is resolved goes down for any reason other than a failure of the interface or of the upstream LSR, routing and LDP goes out of synchronization. This means that the backup ILM remains activated until the next time SPF is run by IGP. By enabling the IGP-LDP synchronization feature, the advertised link metric is changed to the maximum value as soon as the LDP session goes down. This, in turn, triggers an SPF, and LDP downloads a new set of primary and backup ILMs.

The **no** form of this command disables fast upstream switchover for mLDP FECs.

Default

```
no mcast-upstream-frr
```

mp-mbb-time

Syntax

```
mp-mbb-time interval
```

```
no mp-mbb-time
```

Context

```
config>router>ldp
```

Description

This command configures the maximum time a point-to-multipoint transit or bud node must wait before switching over to the new path if the new node does not send an MBB TLV to inform the transit or bud node of the availability of the data plane.

The **no** form of the command sets the wait time to the default.

Default

3 s

Parameters

interval

specifies the MP MBB wait time

Values 1 to 10 seconds

tunnel-down-damp-time

Syntax

tunnel-down-damp-time *seconds*

no tunnel-down-damp-time

Context

config>router>ldp

Description

This command specifies 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 data path; 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), is not be notified immediately. Traffic is still blackholed because the NHLFE has been deprogrammed.

If the FEC gets resolved before the **tunnel-down-damp-time** timer expires, LDP programs the IOM with the new NHLFE and posts a tunnel modify event to 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.

The **no** form of the command resets the damp timer value back to the default value of 3. If the timer value is set to 0, tunnel down events are not dampened but are reported immediately.

Default

3

Parameters

seconds

the time interval that LDP waits before posting a tunnel down event to the TTM

Values 0 to 20

6.12.2.1.3 Interface parameters commands

interface-parameters

Syntax

interface-parameters

Context

config>router>ldp

Description

This command enables the context to configure LDP interfaces and parameters that apply to LDP interfaces.

interface

Syntax

interface *ip-int-namedual-stack*

[no] interface *ip-int-name*

Context

config>router>ldp>if-params

Description

This command enables LDP on the specified IP interface.

The **no** form of the command deletes the LDP interface and all configuration information associated with the LDP interface.

The LDP interface must be disabled using the **shutdown** command before it can be deleted.

You can configure different parameters for IPv4 and IPv6 LDP interfaces by entering **ipv4** or **ipv6** as the next command.

Parameters

dual-stack

distinguishes between configurations created before 7705 SAR Release 9.0 from those created in Release 9.0 or later when the interface node implementation was changed to

include both IPv4 and IPv6 contexts. If the **dual-stack** keyword is used, then the IPv4 interface context is not created. If the keyword is not used, then the IPv4 interface context is automatically created.

When entering an already configured interface, there is no need to provide the keyword; it is ignored if it is included.

By default, all configurations created in Release 9.0 or later include **dual-stack**.

ip-int-name

specifies an existing interface. If the string contains special characters (such as #, \$, spaces), the entire string must be enclosed within double quotes.

ipv4

Syntax

[no] ipv4

Context

```
config>router>ldp>if-params
config>router>ldp>if-params>if
config>router>ldp>targeted-session
```

Description

This command enables the context to configure IPv4 LDP parameters that apply to the interface.

ipv6

Syntax

[no] ipv6

Context

```
config>router>ldp>if-params
config>router>ldp>if-params>if
config>router>ldp>targeted-session
```

Description

This command enables the context to configure IPv6 LDP parameters applied to the interface.

local-lsr-id

Syntax

local-lsr-id {system | interface}

no local-lsr-id**Context**

```
config>router>ldp>if-params>if>ipv4
```

```
config>router>ldp>if-params>if>ipv6
```

Description

This command enables the use of the address of the link LDP interface as the LSR ID to establish an 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 the LSR ID is explicitly configured. This means that targeted LDP (T-LDP) and interface LDP share a common LDP TCP session and therefore a common LDP label space. The system interface must always be configured on the router or the LDP protocol does not come up on the node.

At initial configuration, the LDP session to the peer remains down while the interface is down. If the user changes the LSR ID 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 for the local LSR ID goes down, the LDP session also goes down.

The **interface** option is the recommended setting when static route-LDP synchronization is enabled.

When the **interface** option is selected, the transport connection (TCP) for the link LDP session configured by the [transport-address](#) command is automatically set to **interface**. Having both the **local-lsr-id** and transport address set to the local interface creates two TCP sessions to the peer and therefore two different LDP label spaces: one to the interface IP address for link LDP (L-LDP) and one to the system IP address for T-LDP.

The **no** form of the command resets the **local-lsr-id** to the default value.

Default

system

Parameters**system**

specifies that the system IP address is used to set up the LDP session between peers

interface

specifies that the IP interface address is used to set up the LDP session between peers

p2mp-ipv4**Syntax**

```
p2mp-ipv4 {enable | disable}
```

Context

```
config>router>ldp>session-params>if>ipv4>fec-type-capability
```

```
config>router>ldp>session-params>if>ipv6>fec-type-capability
```

Description

This command enables or disables IPv4 P2MP FEC capability on the interface.

Default

p2mp disable

Parameters

enable | disable

enables or disables IPv4 P2MP FEC capability

p2mp-ipv6

Syntax

p2mp-ipv6 {enable | disable}

Context

config>router>ldp>session-params>if>ipv4>fec-type-capability

config>router>ldp>session-params>if>ipv6>fec-type-capability

Description

This command enables or disables IPv6 P2MP FEC capability on the interface.

Default

p2mp disable

Parameters

enable | disable

enables or disables IPv6 P2MP FEC capability

transport-address

Syntax

transport-address {system | interface}

no transport-address

Context

config>router>ldp>if-params>ipv4

config>router>ldp>if-params>ipv6

config>router>ldp>if-params>if>ipv4

config>router>ldp>if-params>if>ipv6

Description

This command configures the transport address to be used when setting up the LDP TCP sessions. The transport address can be configured globally (applies to all LDP interfaces) or per interface. The most specific value is used.

With the **transport-address** command, you can set up the LDP interface to the connection that 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 address selection situation can also occur when there is 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.

Depending on the first adjacency to be formed, the TCP endpoint is chosen. In other words, if one LDP interface is set up as **transport-address interface** and another as **transport-address system**, then, depending on which adjacency was set up first, the TCP endpoint addresses are determined. After that, because the hello contains the LSR ID, the LDP session can be checked to verify that it is set up and then the adjacency can be matched to the session.

The **no** form of the command:

- at the global level, sets the transport address to the default value
- at the interface level, sets the transport address to the value defined under the global level

Default

system

Parameters

interface

specifies 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

specifies that the system IP address is used to set up the LDP session between neighbors

6.12.2.1.4 Session parameters commands

session-parameters

Syntax

session-parameters

Context

config>router>ldp

Description

This command enables the context to configure peer-specific parameters.

```
peer
```

Syntax

[no] peer *ip-address*

Context

config>router>ldp>session-params

Description

This command configures parameters for an LDP peer.

Default

n/a

Parameters

ip-address

specifies the LDP peer in dotted-decimal notation

```
export-addresses
```

Syntax

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

no export-addresses

Context

config>router>ldp>session-params>peer

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 the command removes the policy from the configuration.

Default

no export-addresses

Parameters

policy-name

the name of an existing export prefix route policy

export-prefixes

Syntax

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

no export-prefixes

Context

config>router>ldp>session-params>peer

Description

This command specifies the export route policy used to determine which prefixes received from other LDP peers are redistributed to this LDP via the LDP session to this peer. Any prefix that has been accepted in the route policy will be exported. Any prefix that has been filtered out of the route policy will not be exported.

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.

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 export policy and targeted sessions.

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

Default

no export-prefixes

Parameters

policy-name

the name of an existing export policy

fec-type-capability

Syntax

fec-type-capability

Context

config>router>ldp>if-params>if>ipv4

config>router>ldp>if-params>if>ipv6

config>router>ldp>session-params>peer

Description

This command enables or disables the advertisement of a FEC type for a specified LDP session or Hello adjacency to a peer.

Default

n/a

p2mp

Syntax

p2mp {enable | disable}

Context

config>router>ldp>session-params>peer>fec-type-capability

Description

This command enables or disables P2MP FEC capability for the session.

Default

p2mp disable

Parameters

enable | disable

enables or disables P2MP FEC capability

prefix-ipv4

Syntax

prefix-ipv4 {enable | disable}

Context

config>router>ldp>if-params>if>ipv4>fec-type-capability

config>router>ldp>if-params>if>ipv6>fec-type-capability

config>router>ldp>session-params>peer>fec-type-capability

Description

This command enables or disables IPv4 prefix FEC capability on the session or interface.

Default

prefix-ipv4 disable

Parameters**enable | disable**

enables or disables IPv4 prefix FEC capability

prefix-ipv6**Syntax****prefix-ipv6 {enable | disable}****Context**

config>router>ldp>if-params>if>ipv4

config>router>ldp>if-params>if>ipv6

config>router>ldp>session-params>peer>fec-type-capability

Description

This command enables or disables IPv6 prefix FEC capability on the session or interface.

Default

prefix-ipv6 disable

Parameters**enable | disable**

enables or disables IPv6 prefix FEC capability

tcp-session-parameters**Syntax****tcp-session-parameters****Context**

config>router>ldp

Description

This command enables the context to configure parameters for the TCP transport session of an LDP session to a remote peer.

Default

n/a

peer-transport

Syntax

[no] peer-transport *ip-address*

Context

config>router>ldp>tcp-session-parameters

Description

This command configures the peer transport address, which is the IPv4 or IPv6 destination address of the TCP connection to the LDP peer.

Default

n/a

Parameters

ip-address

the IPv4 or IPv6 address of the TCP connection to the LDP peer in dotted-decimal notation

auth-keychain

Syntax

auth-keychain *name*

no auth-keychain

Context

config>router>ldp>tcp-session-params>peer-transport

Description

This command associates an authentication keychain with LDP. The keychain is a collection of keys used to authenticate LDP messages from remote peers. The keychain allows the rollover of authentication keys during the lifetime of a session and also supports stronger authentication algorithms than clear text and MD5.

The keychain must already be defined in the **config>system>security>keychain** context.

Either the **authentication-key** command or the **auth-keychain** command can be used by LDP, but both cannot be supported at the same time. If both commands are configured, the **auth-keychain** configuration is applied and the **authentication-key** command is ignored.

By default, authentication is not enabled.

Default

no auth-keychain

Parameters

name

the name of an existing keychain, up to 32 characters

authentication-key

Syntax

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

no authentication-key

Context

config>router>ldp>tcp-session-params>peer-transport

Description

This command specifies the authentication key to be used between LDP peers before establishing sessions. Authentication uses the MD5 message-based digest.

Either the **authentication-key** command or the **auth-keychain** command can be used by LDP, but both cannot be supported at the same time. If both commands are configured, the **auth-keychain** configuration is applied and the **authentication-key** command is ignored.

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

Default

n/a

Parameters

authentication-key

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

hash-key

specifies the hash key. Allowed values are any string up to 33 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, spaces), the entire string must be enclosed within double quotes.

This is useful when a user must configure the parameter; however, for security purposes, the actual unencrypted key value is not provided.

hash

specifies that the key is entered and stored on the node in encrypted form

hash2

specifies that the key is entered and stored on the node in a more complex encrypted form



Note: If neither the **hash** nor **hash2** keyword is specified, the key is entered in clear text. However, for security purposes, the key is stored on the node using hash encryption.

6.12.2.1.5 Targeted session commands

targeted-session

Syntax

targeted-session

Context

config>router>ldp

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.

Default

n/a

disable-targeted-session

Syntax

[no] disable-targeted-session

Context

config>router>ldp>targeted-session

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 the command enables the setup of any targeted sessions.

Default

no disable-targeted-session

peer

Syntax

[no] peer *ip-address*

Context

config>router>ldp>targeted-session

Description

This command configures parameters for an LDP peer.

Default

n/a

Parameters

ip-address

specifies the IPv4 or IPv6 address of the LDP peer

bfd-enable

Syntax

bfd-enable [ipv4][ipv6]

[no] bfd-enable

Context

config>router>ldp>if-params>if

config>router>ldp>targeted-session>peer

Description

This command enables or disables bidirectional forwarding detection (BFD) tracking of the LDP session for the interface or the T-LDP session for the peer.

When BFD is enabled on an LDP interface, the system tracks the next hop of the IPv4 and IPv6 prefixes in addition to tracking the LDP peer address of the Hello adjacency over that link. This is required because LDP can 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 **no** form of the command disables BFD tracking.

Default

n/a

local-lsr-id

Syntax

local-lsr-id *interface-name*

no local-lsr-id

Context

config>router>ldp>targeted-session>peer

Description

This command enables the use of the address of a specific interface as the LSR ID to establish a targeted LDP (T-LDP) adjacency and session with one or more non-directly connected LDP peers. 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, unless the LSR ID is explicitly configured. This means that T-LDP and interface LDP share a common LDP TCP session and therefore a common LDP label space. The system interface must be always be configured on the router or the LDP protocol does not come up on the node.

At initial configuration, the LDP session to the peers remains down while the interface is down. If the user changes the LSR ID 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 for the local LSR ID goes down, the LDP session to all peers using this LSR ID also goes down.

The user-configured LSR ID is used for extended peer discovery to establish the T-LDP hello adjacency. It is also used as the transport address for the LDP TCP 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 the command resets the **local-lsr-id** to the default value, which means that the system interface address is used as the LSR ID.

Default

no local-lsr-id

Parameters

interface-name

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

tunneling

Syntax

[no] tunneling

Context

```
config>router>ldp>targeted-session>peer
```

Description

This command enables LDP over tunnels.

The **no** form of the command disables tunneling.

Default

no tunneling

lsp

Syntax

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

Context

```
config>router>ldp>targeted-session>peer>tunneling
```

Description

This command configures an LSP destined for this peer to be used for tunneling an LDP FEC over RSVP-TE. A maximum of four RSVP-TE LSPs can be used for tunneling LDP FECs to the T-LDP peer.

It is not necessary to specify any RSVP-TE LSP in this context unless there is a need to restrict the tunneling to selected LSPs. All RSVP-TE 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 **config>router>mpls>lsp** *lsp-name* context.

The **no** form of this command removes the LSP association.

Parameters

lsp-name

specifies the name of the LSP

6.12.2.2 Show commands



Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

- [Show router LDP commands](#)
- [Show router LDP bindings commands](#)

6.12.2.2.1 Show router LDP commands

discovery

Syntax

discovery [*state state*] [*detail | summary*] [*adjacency-type type*] [*session ip-addr [label-space]*]

discovery [*state state*] [*detail | summary*] [*adjacency-type type*] [*family*]

discovery interface [*ip-int-name*] [*state state*] [*detail | summary*] [*session ip-addr [label-space]*] [*family*]

discovery peer [*ip-address*] [*state state*] [*detail | summary*] [*session ip-addr [label-space]*]

Context

show>router>ldp

Description

This command displays the status of the interfaces participating in LDP discovery.

Parameters

state

specifies the current operational state of the adjacency

Values established, trying, down

detail

displays detailed information

summary

displays summary information

type

specifies the adjacency type

Values link, targeted

ip-addr

the IP address of the session

label-space

specifies the label space identifier that the router is advertising on the interface

Values 0 to 65535

family

displays either IPv4 or IPv6 LDP session information

ip-int-name

specifies an existing interface. If the string contains special characters (such as #, \$, spaces), the entire string must be enclosed within double quotes.

ip-address
specifies the IP address of the peer

Output

The following outputs are examples of LDP discovery information, and [Table 47: LDP discovery field descriptions](#) describes the fields.

Output example

```
ALU-12# show router ldp discovery
=====
LDP IPv4 Hello Adjacencies
=====
Interface Name          Local Addr      Peer Addr      AdjType State
-----
N/A                     10.10.10.103   10.10.10.93   Targ    Trying
N/A                     10.10.10.103   10.10.10.104  Targ    Estab
to-104                  10.0.0.103     224.0.0.2     Link    Trying
-----
No. of Hello Adjacencies: 3
=====
ALU-12#
```

Output example - show router ldp discovery detail

```
ALU-12# show router ldp discovery detail
=====
LDP IPv4 Hello Adjacencies (Detail)
=====
Peer 10.10.10.93
-----
Local Address      : 10.10.10.103      Peer Address      : 10.10.10.93
Adjacency Type    : Targeted        State             : Trying
-----

Peer 10.10.10.104
-----
Local Address      : 10.10.10.103      Peer Address      : 10.10.10.104
Adjacency Type    : Targeted        State             : Established
Up Time           : 0d 18:26:36      Hold Time Remaining: 38
Hello Mesg Recv   : 76616920        Hello Mesg Sent   : 466580812
Remote Cfg Seq No : 159              Remote IPv4 Address : 198.51.100.255
Local Cfg Seq No  : 1674451         Local IPv4 Address  : 198.51.100.1
-----

Interface "to-104"
-----
Local Address      : 10.0.0.103      Peer Address      : 224.0.0.2
Adjacency Type    : Link            State             : Trying
-----
=====
ALU-12#
```

Table 47: LDP discovery field descriptions

Label	Description
Interface Name	The name of the interface

Label	Description
Local Addr	The IP address of the originating (local) router
Peer Addr	The IP address of the peer
Adj Type	The adjacency type between the LDP peer and LDP session
State	Established – indicates that the adjacency is established
	Trying – indicates that the adjacency is not yet established
No. of Hello Adjacencies	The total number of hello adjacencies discovered
Up Time	The amount of time the adjacency has been enabled
Hold-Time Remaining	The time left before a neighbor is declared to be down
Hello Mesg Recv	The number of Hello messages received for this adjacency
Hello Mesg Sent	The number of Hello messages that have been sent for this adjacency
Remote Cfg Seq No	The configuration sequence number that was in the Hello message received when this adjacency started up. This configuration sequence number changes when there is a change of configuration.
Remote IP Address	The IP address used on the remote end for the LDP session
Local Cfg Seq No	The configuration sequence number that was used in the Hello message sent when this adjacency started up. This configuration sequence number changes when there is a change of configuration.
Local IPv4 Address Local IPv6 Address	The IP address used locally for the LDP session

fec-egress-stats

Syntax

fec-egress-stats [*ip-prefix/mask*]

fec-egress-stats [**active**] [**family**]

Context

show>router>ldp

Description

This command displays LDP FEC egress statistical information.

Parameters

- ip-prefix/mask*

the IP prefix and prefix length associated with the prefix FEC
- active**

displays information from all LDP FECs with statistics collection enabled
- family**

displays either IPv4 or IPv6 LDP information
- Values**

ipv4 or ipv6

Output

The following output is an example of LDP FEC egress statistical information.

Output example

```
e.show>router>ldp>bindings>fec-egress-stats 2.2.2.2/32
=====
LDP Egress Statistics
=====
-----
FEC Prefix/Mask      : 10.10.10.29/32
-----
Collect Stats       : Disabled           Accounting Plcy.    : None
Admin State        : Up
FC BE
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets       : 0                   OutProf Octets     : 0
FC L2
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets       : 0                   OutProf Octets     : 0
FC AF
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets       : 0                   OutProf Octets     : 0
FC L1
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets       : 0                   OutProf Octets     : 0
FC H2
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets       : 0                   OutProf Octets     : 0
FC EF
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets       : 0                   OutProf Octets     : 0
FC H1
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets       : 0                   OutProf Octets     : 0
FC NC
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets       : 0                   OutProf Octets     : 0
-----
LDP Egress Statistics : 1
=====
```

fec-originate

Syntax

fec-originate [*ip-address/mask*] [**operation-type**]
fec-originate [**operation-type**] [*family*]

Context

show>router>ldp

Description

This command displays LDP static prefix FECs.

Parameters

- ip-address/mask*

specifies the IPv4 or IPv6 address
- operation-type**

specifies the operation type to display

Values pop | swap
- family*

the address family filter

Values ipv4 or ipv6

Output

The following output is an example of FEC originate information, and [Table 48: FEC-originate field descriptions](#) describes the fields.

Output example

```
*A:ALU-12# show router ldp fec-originate
```

LDP Static Prefix FECs					
Prefix	NHType	NextHop	IngLabel	EgrLabel	OperIngLabel
10.1.0.0/16	Pop	n/a	--	--	0
10.1.0.1/32	Pop	n/a	--	--	0
10.1.0.2/32	Pop	n/a	--	--	0
10.1.0.3/32	Pop	n/a	--	--	0
10.1.0.4/32	Pop	n/a	--	--	0
10.1.0.5/32	Pop	n/a	--	--	0
10.1.0.6/32	Pop	n/a	--	--	0
10.1.0.7/32	Pop	n/a	--	--	0
10.1.0.8/32	Pop	n/a	--	--	0
10.1.0.9/32	Pop	n/a	--	--	0
...					
10.251.0.0/16	Pop	n/a	--	--	0
10.252.0.0/16	Pop	n/a	--	--	0

10.253.0.0/16	Pop	n/a	--	--	0
10.254.0.0/16	Pop	n/a	--	--	0

No. of FECs: 508					
=====					
LDP IPv6 Static Prefix FECs					
=====					
Prefix	NHType	NextHop	IngLabel	EgrLabel	OperIngLabel

No Matching Entries Found					
=====					
*A:ALU-12#					

Table 48: FEC-originate field descriptions

Label	Description
Prefix	The static prefix FEC
NHType	The type of next hop for this entry: Unknown: the next-hop type has not been set IP Addr: the next hop is an IP address Pop: there is no next hop; label is popped and packet routed Unnumbered: the next hop is an unnumbered interface
Next Hop	The IP address of the next hop, or Unnumbered for unnumbered interfaces
IngLabel	The label that is advertised to the upstream peer. If this variable is set to the default value of 4294967295, the ingress label is dynamically assigned by the label manager.
EgrLabel	The egress label associated with this next-hop entry. The LSR swaps the incoming label with the configured egress label. If this egress label has a value of 4294967295, the LSR pops the incoming label.
OperIngLabel	The actual or operational value of the label that was advertised to the upstream peer

interface

Syntax

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

Context

show>router>ldp

Description

This command displays configuration information about LDP interfaces.

Parameters

- ip-int-name*
specifies an existing interface. If the string contains special characters (such as #, \$, spaces), the entire string must be enclosed within double quotes.
- ip-address*
identifies the LDP neighbor by IP address
- detail**
displays detailed information

Output

The following output is an example of LDP interface information, and [Table 49: LDP interface field descriptions](#) describes the fields.

Output example

```
A:ALU-12# show router ldp interface
=====
LDP Interfaces
=====
Interface          Adm/Opr  Hello  Hold  KA   KA   Transport
Sub-Interface<s>  Adm/Opr  Fctr   Time Fctr  Time Address
-----
ip-10.10.1.1      Up/Up
  ipv4            Up/Dwn  3      15   3    30   System
  ipv6            Up/Dwn  3      15   3    30   System
-----
No. of Interfaces: 1
=====
A:ALU-12#

A:ALU-12>show>router>ldp# interface detail
=====
LDP Interfaces (Detail)
=====
Interface "back"
-----
BASE
-----
Admin State      : Up           Oper State      : Down
BFDD Status      : ipv4
-----
IPv4
-----
IPv4 Admin State:  Up           IPv4 Oper State : Down
IPv4 Oper Down Rea*: interfaceDown
Hold Time       : 1000          Hello Factor    : 15
Keepalive Timeout : 1000        Keepalive Factor : 15
Transport Addr   : System       Last Modified   : 08/08/2007 09:50:15
Active Adjacencies : 0
Tunneling        : Disabled
Lsp Name         : None
```

```
=====
A:ALU-12>show>router>ldp#
```

Table 49: LDP interface field descriptions

Label	Description
Interface	The interface associated with the LDP instance
Adm	Up – indicates that the LDP is administratively enabled
	Down – indicates that the LDP is administratively disabled
Opr	Up – indicates that the LDP is operationally enabled
	Down – indicates that the LDP is operationally disabled
Hello Fctr	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	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.
KA Fctr	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.
KA Time	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).
Transport Address	The transport address entity
No. of Interfaces	The total number of LDP interfaces
Oper Down Reason	The reason for the LSP being in the down state
Active Adjacencies	The number of active adjacencies
Last Modified	The time of the last modification to the LDP interface
Lsp Name	The LSP name

parameters

Syntax

parameters

Context

show>router>ldp

Description

This command displays configuration information about LDP parameters.

Output

The following output is an example of LDP parameters information, and [Table 50: LDP parameters field descriptions](#) describes the fields.

Output example

```
A:ALU-12# show router ldp parameters
=====
LDP Parameters (IPv4 LSR ID 10.20.1.1:0)
                (IPv6 LSR ID 3ffe::a14:101[0])
=====
-----
Graceful Restart Parameters
-----
Graceful Restart      : Disabled
Nbor Liveness Time   : 120 sec           Max Recovery Time   : 120
-----
IPv4 Interface Parameters
-----
Keepalive Timeout    : 30 sec           Keepalive Factor      : 3
Hold Time            : 15 sec           Hello Factor           : 3
Transport Address    : system
-----
IPv6 Interface Parameters
-----
Keepalive Timeout    : 30 sec           Keepalive Factor      : 3
Hold Time            : 15 sec           Hello Factor           : 3
Transport Address    : system
-----
Targeted Session Parameters
-----
Import Pfx Policies: None              Export Pfx Policies : None
Prefer Tunl-in-Tunl: Disabled          SDP Auto Targ Sess  : Enabled
-----
IPv4 Targeted Session Parameters
-----
Keepalive Timeout    : 30 sec           Keepalive Factor      : 3
Hold Time            : 15 sec           Hello Factor           : 3
Hello Reduction      : Disabled         Hello Reduction Fctr: 3
-----
IPv6 Targeted Session Parameters
-----
Keepalive Timeout    : 40 sec           Keepalive Factor      : 4
Hold Time            : 45 sec           Hello Factor           : 3
Hello Reduction      : Disabled         Hello Reduction Fctr: 3
```


=====

A:ALU-12A#

Table 50: LDP parameters field descriptions

Label	Description
Graceful Restart Parameters	
Nbor Liveliness Time	The neighbor liveness time
Max Recovery Time	The local maximum recovery time
IPv4 Interface Parameters	
IPv6 Interface Parameters	
Keepalive Timeout	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).
Keepalive Factor	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.
Hold Time	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.
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.
Propagate Policy	Specifies whether the LSR should generate FECs and which FECs it should generate
	system – indicates that the LDP distributes label bindings only for the router's system IP address
	interface – indicates that the LDP distributes label bindings for all LDP interfaces
	all – indicates that the LDP distributes label bindings for all prefixes in the routing table
	none – indicates that the LDP does not distribute any label bindings

Label	Description
Transport Address	interface – the interface IP address is used to set up the LDP session between neighbors. If multiple interfaces exist between two neighbors, the interface mode cannot be used because only one LDP session is actually set up between the two neighbors.
	system – the system IP address is used to set up the LDP session between neighbors
Label-Distribution	The label distribution method
Label-Retention	liberal – all advertised label mappings are retained whether they are from a valid next hop or not. When the label distribution value is downstream unsolicited, a router may receive label bindings for the same destination for all its neighbors. Labels for the non-next-hops for the FECs are retained in the software but not used. When a network topology change occurs where a non-next-hop becomes a true next hop, the label received earlier is then used.
	conservative – advertised label mappings are retained only if they are used to forward packets; for example if the label came from a valid next hop. Label bindings received from non-next-hops for each FEC are discarded.
Control Mode	ordered – label bindings are not distributed in response to a label request until a label binding has been received from the next hop for the destination
	independent – label bindings are distributed immediately in response to a label request even if a label binding has not yet been received from the next hop for the destination
Route Preference	The route preference assigned to LDP routes. When multiple routes are available to a destination, the route with the lowest preference is used. This value is only applicable to LDP interfaces and not for targeted sessions.
IPv4 Targeted Session Parameters	
IPv6 Targeted Session Parameters	
Keepalive Timeout	The factor used to derive the keepalive interval
Keepalive Factor	The time interval, in seconds, that LDP waits before tearing down the session
Hold Time	The time left before a neighbor is declared to be down
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.

Label	Description
	Disable – indicates that no authentication is being used
Passive Mode	True – indicates that LDP responds only when it gets a connect request from a peer and does not attempt to actively connect to its neighbors
	False – indicates that LDP actively tries to connect to its peers
Targeted Sessions	Enabled – indicates that targeted sessions are enabled
	Disabled – indicates that targeted sessions are disabled

session

Syntax

```

session [ip-addr [:label-space]] local-addresses [sent | recv] ip-addr ip-address
session [ip-addr [:label-space]] [session-type] [state state] [detail | summary]
session [ip-addr [:label-space]] local-addresses [sent | recv] [family]
session [ip-addr [:label-space]] statistics [packet-type] [session-type]
session statistics [packet-type] [session-type] [family]
session [session-type] [state state] [detail | summary] [family]

```

Context

```
show>router>ldp
```

Description

This command displays configuration information about LDP sessions.

Parameters

ip-addr

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

summary

displays summary information

ip-address

specifies the IP address

- state

specifies the current operational state of the adjacency

Valuesestablished, trying, down
- packet-type

specifies the packet type

Valueshello, keepalive, init, label, notification, address
- family

displays either IPv4 or IPv6 LDP session information
- session-type

specifies the session type

Valueslink, targeted, both

Output

The following output is an example of LDP session information, and [Table 51: LDP session field descriptions](#) describes the fields.

Output example

```
ALU-12# show router ldp session
=====
LDP IPv4 Sessions
=====
Peer LDP Id      Adj Type State      Msg Sent  Msg Recv  Up Time
-----
10.10.10.104:0   Targeted Established 13943     13947     0d 21:12:41
-----
No. of IPv4 Sessions: 1

=====
LDP IPv6 Sessions
=====
Peer LDP Id      Adj Type      State      Msg Sent  Msg Recv  Up Time
-----
3ffe::a14:102[0] Link          Established 1788      1792      0d 01:19:19
3ffe::a14:103[0] Link          Established 1789      1788      0d 01:19:19
-----
No. of IPv6 Sessions: 2

=====
ALU-12#

ALU-12# show router ldp session detail
=====
LDP IPv4 Sessions (Detail)
=====
Session with Peer 10.1.1.33:0
-----
Adjacency Type   : Link          State           : Established
Up Time          : 0d 00:03:51
Max PDU Length   : 4096          KA/Hold Time Remaining: 26
```

```

Link Adjacencies : 1          Targeted Adjacencies : 0
Local Address    : 10.1.1.30  Peer Address     : 10.1.1.33
Local TCP Port   : 646        Peer TCP Port    : 50232
Local KA Timeout : 30         Peer KA Timeout  : 30
Mesg Sent        : 89         Mesg Recv        : 126
FECs Sent        : 3          FECs Recv        : 3
GR State         : Not Capable
Nbr Liveness Time : 0         Max Recovery Time  : 0
Number of Restart : 0         Last Restart Time  : Never
Advertise        : Address
-----
Session with Peer 10.1.1.57:0
-----
Adjacency Type   : Targeted    State                : Established
Up Time          : 0d 00:03:49
Max PDU Length   : 4096        KA/Hold Time Remaining: 36
Link Adjacencies : 0          Targeted Adjacencies : 1
Local Address    : 10.1.1.30  Peer Address     : 10.1.1.57
Local TCP Port   : 646        Peer TCP Port    : 49574
Local KA Timeout : 40         Peer KA Timeout  : 40
Mesg Sent        : 55         Mesg Recv        : 61
FECs Sent        : 11         FECs Recv        : 8
GR State         : Not Capable
Nbr Liveness Time : 0         Max Recovery Time  : 0
Number of Restart : 0         Last Restart Time  : Never
Advertise        : Address/Servi*
=====
ALU-12#

```

Table 51: LDP session field descriptions

Label	Description
Peer LDP Id	The IP address of the LDP peer
Adj Type	The adjacency type between the LDP peer and LDP session that 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
Msg Sent	The number of messages sent
Msg Rcvd	The number of messages received
Up Time	The amount of time the adjacency has been enabled

session-parameters

Syntax

session-parameters [*family*]
session-parameters [*peer-ip-address*]

Context

show>router>ldp

Description

This command displays LDP peer information.

Parameters

- family*

the address family filter

Values ipv4 or ipv6
- peer-ip-address*

specifies the peer IP address

Output

The following output is an example of LDP session-parameters information.

Output example

```
A:ALU-12# show router ldp session-parameters
=====
LDP IPv4 Session Parameters
=====
-----
Peer : 10.2.3.4
-----
DOD                : Disabled          Adv Adj Addr Only : Disabled
FEC129             : 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: Enabled
P2MP Fec Cap       : Enabled
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
=====
```

status

Syntax
status

Context
show>router>ldp

Description
This command displays LDP status information.

Output
The following output is an example of LDP status information, and [Table 52: LDP status field descriptions](#) describes the fields.

Output example

```
A:ALU-12# show router ldp status

=====
LDP Status for IPv4 LSR ID 2.2.2.2
                IPv6 LSR ID ::
=====
Created at      : 10/18/17 13:08:17
Last Change    : 10/18/17 13:08:17
Admin State    : Up
IPv4 Oper State : Up
IPv4 Up Time   : 1d 01:35:18
IPv4 Oper Down Rea*: n/a
IPv4 Oper Down Eve*: 0
Tunn Down Damp Time: 3 sec
Label Withdraw Del*: 0 sec
Short. TTL Local  : Enabled
ConsiderSysIPByPol*: Disabled
Import Policies  : None
Tunl Exp Policies : None
FRR              : Disabled
Mcast Upst ASBR FRR: Disabled
MP MBB Time     : 3
Aggregate Prefix : False
Class Forwarding : Disabled
Legacy LSR Interop : False
Entropy Label Capa*: False
Generate Basic FEC : Disabled
Resolve Via Mcast *: Disabled

-----
Capabilities
-----
Dynamic          : Enabled
IPv4 Prefix Fec  : Enabled
Service Fec128   : Enabled
MP MBB           : Enabled
Unrecognized Notif*: Enabled
P2MP             : Enabled
IPv6 Prefix Fec  : Enabled
Service Fec129   : Enabled
Overload         : Enabled

=====
* indicates that the corresponding row element may have been truncated.
```

Table 52: LDP status field descriptions

Label	Description
Created at	The date and time that the LDP instance was created
Last Change	The date and time that the LDP instance was last modified
Admin State	Up – indicates that LDP is administratively enabled
	Down – indicates that LDP is administratively disabled
IPv4 Oper State	Up – indicates that LDP is operationally enabled
IPv6 Oper State	Down – indicates that LDP is operationally disabled
IPv4 Up Time	The time, in hundredths of seconds, that the LDP instance has been operationally up
IPv6 Up Time	
IPv4 Oper Down Time	The time, in hundredths of seconds, that the LDP instance has been operationally down
IPv6 Oper Down Time	
IPv4 Oper Down Events	The number of times the LDP instance has gone operationally down since the instance was created
IPv6 Oper Down Events	
Import Policies	The import policy associated with the LDP instance
Active Adjacencies	The number of active adjacencies (established sessions) associated with the LDP instance
Active Sessions	The number of active sessions (session in some form of creation) associated with the LDP instance
Active Interfaces	The number of active (operationally up) interfaces associated with the LDP instance
Inactive Interfaces	The number of inactive (operationally down) interfaces associated with the LDP instance
Active Peers	The number of active LDP peers
Inactive Peers	The number of inactive LDP peers
Addr FECs Sent	The number of labels that have been sent to the peer associated with this FEC
Addr FECs Recv	The number of labels that have been received from the peer associated with this FEC

Label	Description
Serv FECs Sent	The number of labels that have been sent to the peer associated with this FEC
Serv FECs Recv	The number of labels that have been received from the peer associated with this FEC
Attempted Sessions	The total number of attempted sessions for this LDP instance
No Hello Err	The total number of "Session Rejected" or "No Hello Error" notification messages sent or received by this LDP instance
Param Adv Err	The total number of "Session Rejected" or "Parameters Advertisement Mode Error" notification messages sent or received by this LDP instance
Max PDU Err	The total number of "Session Rejected" or "Parameters Max PDU Length Error" notification messages sent or received by this LDP instance
Label Range Err	The total number of "Session Rejected" or "Parameters Label Range Error" notification messages sent or received by this LDP instance
Bad LDP Id Err	The number of bad LDP identifier fatal errors detected for sessions associated with this LDP instance
Bad PDU Len Err	The number of bad PDU length fatal errors detected for sessions associated with this LDP instance
Bad Mesg Len Err	The number of bad message length fatal errors detected for sessions associated with this LDP instance
Bad TLV Len Err	The number of bad TLV length fatal errors detected for sessions associated with this LDP instance
Malformed TLV Err	The number of malformed TLV value fatal errors detected for sessions associated with this LDP instance
Keepalive Expired Err	The number of session keepalive timer expired errors detected for sessions associated with this LDP instance
Shutdown Notif Sent	The number of shutdown notifications sent related to sessions associated with this LDP instance
Shutdown Notif Recv	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]
```

Context

```
show>router>ldp
```

Description

This command displays configuration information about LDP targeted peers.

Parameters

- ip-address*
specifies the IP address of the LDP peer
- detail**
displays detailed information
- family*
the address family filter
 - Values** ipv4 or ipv6

Output

The following output is an example of LDP peer information, and [Table 53: LDP targeted peer field descriptions](#) describes the fields.

Output example

A:ALU-12# show router ldp targ-peer

LDP IPv4 Targeted Peers							
Peer	Adm/ Opr	Hello Fctr	Hold Time	KA Fctr	KA Time	Passive Mode	Auto Created
10.10.10.93	Up/Up	3	45	4	40	Disabled	Yes
10.10.10.104	Up/Up	3	45	4	40	Disabled	Yes
No. of IPv4 Targeted Peers: 2							
LDP IPv6 Targeted Peers							
Peer	Adm/ Opr	Hello Fctr	Hold Time	KA Fctr	KA Time	Auto Created	
3ffe::a0a:203	Up/Up	3	15	3	30	no	

```

No. of IPv6 Targeted Peers: 1
=====
A:ALU-12#

A:ALU-12# show router ldp targ-peer detail
=====
LDP IPv4 Targeted Peers (Detail)
=====
-----
Peer 10.2.3.4
-----
Admin State       : Up           Oper State       : Down
Hold Time        : 45           Hello Factor     : 3
Keepalive Timeout : 40          Keepalive Factor : 4
Passive Mode      : Disabled     Last Modified    : 05/01/2008 21:44:17
Active Adjacencies : 0           Auto Created     : No
Tunneling         : None
Lsp Name          : None
=====
No. of IPv4 Targeted Peers: 1
=====

=====
LDP IPv6 Targeted Peers
=====
No Matching Entries Found
=====
A:ALU-12#

```

Table 53: LDP targeted peer field descriptions

Label	Description
Peer	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	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	The value by which the keepalive timeout should be divided to give the keepalive time; that is, the time interval, in seconds,

Label	Description
	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	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).
Passive Mode	The mode used to set up LDP sessions. This value is only applicable to targeted sessions and not to LDP interfaces. This mode is always set to False.
	True – indicates that LDP responds only when it gets a connect request from a peer and does not attempt to actively connect to its neighbors
	False – indicates that LDP actively tries to connect to its peers
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	The total number of LDP peers
LSP	The LSP name

tcp-session-parameters

Syntax

tcp-session-parameters [*family*]

tcp-session-parameters [**keychain** *keychain*]

tcp-session-parameters *transport-peer-ip-address*

Context

show>router>ldp

Description

This command displays information about the TCP transport session of an LDP peer.

Parameters

family

displays either IPv4 or IPv6 LDP session information

keychain

specifies the authentication keychain name up to 32 characters in length

Values ipv4 or ipv6

transport-peer-ip-address

specifies the IP address of the transport peer

Output

The following output is an example of TCP session parameter information.

Output example

```
*A:Dut-A# show router ldp tcp-session-parameters
=====
LDP IPv4 TCP Session Parameters
=====
-----
Peer Transport: 10.20.1.2
-----
Authentication Key : Disabled Path MTU Discovery : Disabled
Auth key chain : LdpAuth Min-TTL : 0
-----
Peer Transport: 10.20.1.3
-----
Authentication Key : Disabled Path MTU Discovery : Disabled
Auth key chain : LdpAuth Min-TTL : 0
=====
No. of IPv4 Peers: 2
=====
LDP IPv6 TCP Session Parameters
=====
-----
Peer Transport: 3ffe::a14:102
-----
Authentication Key : Disabled Path MTU Discovery : Disabled
Auth key chain : LdpAuth Min-TTL : 0
-----
Peer Transport: 3ffe::a14:103
-----
Authentication Key : Disabled Path MTU Discovery : Disabled
Auth key chain : LdpAuth Min-TTL : 0
=====
No. of IPv6 Peers: 2
=====
```

6.12.2.2.2 Show router LDP bindings commands**bindings****Syntax**

bindings

Context

show>router>ldp

Description

This command shows LDP bindings information. The **bindings** command can be used with the following keywords:

- **active**: displays LDP active bindings
- **detail**: displays details of LDP bindings
- **ipv4**: displays LDP IPv4 bindings
- **ipv6**: displays LDP IPv6 bindings
- **label-type**: displays LDP FEC bindings by matching labels
- **p2mp**: displays LDP P2MP FEC bindings
- **prefixes**: displays LDP prefix FEC bindings
- **services**: displays LDP service FEC bindings
- **session**: displays LDP FEC bindings by matching peer LSR ID
- **summary**: displays summary of LDP bindings

Output

The following output is an example of LDP bindings parameter information, and [Table 54: LDP bindings field descriptions](#) describes the fields.

Output example

```
*A:Sa18 Dut-B>show>router>ldp# bindings
=====
LDP Bindings (IPv4 LSR ID 2.2.2.2)
(IPv6 LSR 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
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value),
        H - Hpipe Service
        LF - Lower FEC, UF - Upper FEC
=====
LDP IPv4 Prefix Bindings
=====
Prefix                                     IngLbl                                     EgrLbl
Peer                                     EgrIntf/LspId
EgrNextHop
-----
10.1.1.1/32                               --                                     131071
10.1.1.1:0                               1/1/1:100
10.1.1.1
10.2.2.2/32                               131071U
10.1.1.1:0                               --
--
-----
No. of IPv4 Prefix Bindings: 2
=====
```

```

=====
LDP IPv6 Prefix Bindings
=====
Prefix                                IngLbl                                EgrLbl
Peer                                EgrIntf/LspId
EgrNextHop
-----
3ffe::a14:101/128                    262142U                               --
3ffe::a14:102[0]                      --
-----
No. of IPv6 Prefix Bindings: 2
=====

LDP Generic IPv4 P2MP Bindings
=====
P2MP-Id
RootAddr                                Interface                                IngLbl    EgrLbl
EgrNH                                EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

LDP Generic IPv6 P2MP Bindings
=====
P2MP-Id
RootAddr                                Interface                                IngLbl    EgrLbl
EgrNH                                EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

LDP In-Band-SSM IPv4 P2MP Bindings
=====
Source
Group                                Interface                                IngLbl    EgrLbl
RootAddr                                EgrIf/LspId
EgrNH
Peer
-----
No Matching Entries Found
=====

LDP In-Band-SSM IPv6 P2MP Bindings
=====
Source
Group                                Interface                                IngLbl    EgrLbl
RootAddr                                EgrIf/LspId
EgrNH
Peer
-----
No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv4 P2MP Bindings
=====
Source
Group                                RD                                Interface                                IngLbl    EgrLbl
RootAddr                                EgrIf/LspId
EgrNH
Peer
-----
No Matching Entries Found
=====

```

```

=====
LDP In-Band-VPN-SSM IPv6 P2MP Bindings
=====
Source
Group                                RD
RootAddr                            Interface    IngLbl    EgrLbl
EgrNH                               EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

LDP Recursive with In-Band-SSM IPv4 P2MP Bindings
=====
RootAddr
InnerRootAddr
Source
Group                                Interface    IngLbl    EgrLbl
EgrNH                               EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

LDP Recursive with In-Band-SSM IPv6 P2MP Bindings
=====
RootAddr
InnerRootAddr
Source
Group                                Interface    IngLbl    EgrLbl
EgrNH                               EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

LDP VPN Recursive with Generic IPv4 P2MP Bindings
=====
P2MP-Id                                RD
RootAddr                            Interface
InnerRootAddr                        IngLbl    EgrLbl
EgrNH                               EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

LDP VPN Recursive with Generic IPv6 P2MP Bindings
=====
P2MP-Id                                RD
RootAddr                            Interface
InnerRootAddr                        IngLbl    EgrLbl
EgrNH                               EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

LDP GRT Recursive with Generic IPv4 P2MP Bindings
=====
P2MP-Id
RootAddr                            Interface    IngLbl    EgrLbl
InnerRootAddr

```



```

EgrNH                               EgrIf/LspId
Peer
-----
No Matching Entries Found
=====
LDP GRT Recursive with Generic IPv6 P2MP Bindings
=====
P2MP-Id
RootAddr                          Interface      IngLbl    EgrLbl
InnerRootAddr                     EgrIf/LspId
EgrNH
Peer
-----
No Matching Entries Found
=====
LDP Service FEC 128 Bindings
=====
Type                               VCId      SDPIId    IngLbl    LMTU
Peer                               SvcId     EgrLbl    EgrLbl    RMTU
-----
?-Eth                             10        R. Src    --        None
10.1.1.1:0                        Ukwn      131070D  1500
-----
No. of VC Labels: 1
=====
LDP Service FEC 129 Bindings
=====
SAII                               AGII      IngLbl    LMTU
TAII                               Type      EgrLbl    RMTU
Peer                               SvcId     SDPIId
-----
No Matching Entries Found
=====
*A:Sar18 Dut-B>show>router>ldp#

```

Table 54: LDP bindings field descriptions

Label	Description	
Legend	U: Label In Use	I: IES Service
	N: Label Not In Use	R: VPRN service
	W: Label Withdrawn	P: Ipipe Service
	S: Status Signaled Up	WP: Label Withdraw Pending
	D: Status Signaled Down	C: Cpipe Service
	E: Epipe Service	BU: Alternate for Fast Re-Route
	V: VPLS Service	TLV: (Type, Length: Value)
	M: Mirror Service	H: Hpipe
	A: Apipe Service	LF: Lower FEC
	F: Fpipe Service	UF: Upper FEC

Label	Description
Type	The service type exchanging labels in the SDP. The possible types displayed are Epipe, Spoke, and Unknown.
VCId	The value used by each end of an SDP tunnel to identify the VC
SvcID	Identifies the service in the service domain
SDPId	Identifies the SDP in the service domain
Peer	The IP address of the peer
IngLbl	The ingress LDP label
	U – indicates that the label is in use
	R – indicates that the label has been released
EgrLbl	The egress LDP label
LMTU	The local MTU value
RMTU	The remote MTU value
No. of Prefix Bindings	The total number of LDP bindings on the router
EgrIntf/LspId	The egress interface LSP ID
EgrNextHop	The egress next-hop address, or Unnumbered for unnumbered interfaces
AGI Type	The address group identifier (AGI)
SAII Peer	The source attachment individual identifier (SAII)
TAII EgrLbl	The target attachment individual identifier (TAII)
Prefix	The FEC address
P2MP-Id	The Internal identifier of Point to Multi-point LSP
RootAddr	The Root Address (only IPV4)
Interface	The Logical Interface ID
EgrNH	The egress next-hop address
Source	The Source Address
Group	The Multicast Group address
InnerRootAddr	The Inner Root Address
No. of VC Labels	The total number of VC labels

Label	Description
No. of Service Bindings	The total number of service bindings
Vc-switching	Not applicable – always indicates No
Egr. Flags	Specifies the egress flags, if any
Egr. Ctl Word	Indicates whether egress control words are used
Egr. Status Bits	Indicates whether egress status bits are supported
Egr If Name	The egress interface name
Ing. Flags	Specifies the ingress flags, if any
Ing. Ctl Word	Indicates whether ingress control words are used
Ing. Status Bits	Indicates whether ingress status bits are supported
Metric	The metric of the LSP
Mtu	The size of the MTU for the global FEC or tunnel to which the LDP binding is applied
Op	The operation performed on the ingress or egress label in the LDP stack (push or pop)

active

Syntax

active detail [*family*][**egress-if** *port-id*]

active detail [*family*][**egress-lsp** *tunnel-id*]

active detail [**egress-nh** *ip-address*] [*family*]

active egress-if *port-id* [**summary** | **detail**] [*family*]

active egress-lsp *tunnel-id*[**summary** | **detail**] [*family*]

active egress-nh [*family*] [**summary** | **detail**] *ip-address*

active ipv4 [**summary** | **detail**][**egress-if** *port-id*]

active ipv4 [**summary** | **detail**][**egress-lsp** *tunnel-id*]

active ipv4 [**summary** | **detail**][**egress-nh** *ip-address*]

active ipv6 [**summary** | **detail**] [**egress-if** *port-id*]

active ipv6 [**summary** | **detail**][**egress-nh** *ip-address*]

active ipv6 [**summary** | **detail**][**egress-lsp** *tunnel-id*]

active p2mp p2mp-id *identifier* *root ip-address* [**summary** | **detail**] [**egress-if** *port-id*]

active p2mp p2mp-id *identifier* *root ip-address* [**summary** | **detail**] [**egress-lsp** *tunnel-id*]

active p2mp p2mp-id *identifier* **root** *ip-address* [**summary** | **detail**] [**egress-nh** *ip-address*]
active p2mp [*family*] [**summary** | **detail**] [**egress-if** *port-id*] [**opaque-type** *opaque-type*]
active p2mp [*family*] [**summary** | **detail**] [**egress-lsp** *tunnel-id*] [**opaque-type** *opaque-type*]
active p2mp [*family*] [**summary** | **detail**] [**egress-nh** *ip-address*] [**opaque-type** *opaque-type*]
active p2mp source *ip-address* **group** *mcast-address* **root** *ip-address* [**summary** | **detail**] [**egress-if** *port-id*] **inner-root** *ip-address*
active p2mp source *ip-address* **group** *mcast-address* **root** *ip-address* [**summary** | **detail**] [**egress-lsp** *tunnel-id*] **inner-root** *ip-address*
active p2mp source *ip-address* **group** *mcast-address* **root** *ip-address* [**summary** | **detail**] [**egress-nh** *ip-address*] **inner-root** *ip-address*
active p2mp source *ip-address* **group** *mcast-address* **root** *ip-address* [**rd** *rd*] [**summary** | **detail**] [**egress-if** *port-id*]
active p2mp source *ip-address* **group** *mcast-address* **root** *ip-address* [**rd** *rd*] [**summary** | **detail**] [**egress-lsp** *tunnel-id*]
active p2mp source *ip-address* **group** *mcast-address* **root** *ip-address* [**rd** *rd*] [**summary** | **detail**] [**egress-nh** *ip-address*]
active p2mp source *ip-address* **group** *mcast-address* *family* [**rd** *rd*] [**summary** | **detail**] [**innermost-root** *ip-address*]
active prefixes [*family*] [**summary** | **detail**] [**egress-if** *port-id*]
active prefixes [*family*] [**summary** | **detail**] [**egress-lsp** *tunnel-id*]
active prefixes [**egress-nh** *ip-address*] [*family*] [**summary** | **detail**]
active prefixes prefix *ip-prefix/ip-prefix-length* [**summary** | **detail**] [**egress-if** *port-id*]
active prefixes prefix *ip-prefix/ip-prefix-length* [**summary** | **detail**] [**egress-lsp** *tunnel-id*]
active prefixes prefix *ip-prefix/ip-prefix-length* [**egress-nh** *ip-address*] [**summary** | **detail**]
active summary [*family*] [**egress-if** *port-id*]
active summary [*family*] [**egress-lsp** *tunnel-id*]
active summary [**egress-nh** *ip-address*] [*family*]

Context

show>router>ldp>bindings

Description

This command displays information about LDP active bindings.

Parameters

detail

displays detailed information

summary

displays information in a summarized format

family

displays either IPv4 or IPv6 active LDP information

Values ipv4 or ipv6

egress-if *port-id*

displays LDP active bindings by matching **egress-if**

egress-lsp *tunnel-id*

displays LDP active bindings by matching on the egress RSVP-TE LSP *tunnel-id* for LDP FECs that are tunneled over an RSVP-TE LSP. The *tunnel-id* for the RSVP-TE LSP can be found in the output of the **show router mpls lsp detail** command and in the **show router tunnel-table** command. It is not the **Path LSP ID** shown in the output of the **show router mpls lsp path detail** command.

Values 0 to 4294967295

egress-nh *ip-address*

displays LDP active bindings by matching **egress-nh**

opaque-type *opaque-type*

specifies the type of a multipoint opaque value element

Values generic, ssm, vpn-ssm, recursive-ssm, vpn-recursive, grt-recursive

inner-root *ip-address*

displays recursive FECs whose inner root address matches the specified address

innermost-root *ip-address*

displays recursive FECs whose inner root address matches the specified address and non-recursive FECs that have a root address that matches the specified address

p2mp source *ip-address*

displays LDP active P2MP source bindings

p2mp-id *identifier*

displays LDP active P2MP identifier bindings

Values 0 to 4294967295

group *mcast-address*

displays the P2MP group multicast address bindings

Values ipv4-mcast-addr or ipv6-mcast-addr

root *ip-address*

displays root IP address information

rd *rd*

displays information for the route distinguisher

Values *rd* : *ip-addr:comm-val* or *2byte-asnumber:ext-comm-val* or *4byte-asnumber:comm-val*

prefix ip-prefix/ip-prefix-length

displays information for the specified IP prefix and mask length

Output

The following output is an example of LDP active bindings information, and [Table 54: LDP bindings field descriptions](#) describes the fields.

Output example

```
*A:Sa18 Dut-B>show>router>ldp# bindings active
=====
LDP Bindings (IPv4 LSR ID 10.20.1.3:0)
      (IPv6 LSR ID 3ffe::a14:103[0])
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
      WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
      LF - Lower FEC, UF - Upper FEC
      (S) - Static          (M) - Multi-homed Secondary Support
      (B) - BGP Next Hop    (BU) - Alternate Next-hop for Fast Re-Route
      (I) - SR-ISIS Next Hop (O) - SR-OSPF Next Hop
=====
LDP IPv4 Prefix Bindings (Active)
=====
Prefix          Op          IngLbl  EgrLbl
EgrNextHop      EgrIf/LspId
-----
10.20.1.1/32    Push        --      262143
10.10.2.1       1/1/1
10.20.1.1/32    Swap        262141  262143
10.10.2.1       1/1/1
-----
No. of IPv4 Prefix Active Bindings: 2
=====
LDP IPv6 Prefix Bindings (Active)
=====
Prefix          Op          IngLbl  EgrLbl
EgrNextHop      EgrIf/LspId
-----
3ffe::a14:101/128 Push        --      262142
fe80::21        1/1/1
3ffe::a14:101/128 Swap        262136  262142
fe80::21        1/1/1
-----
No. of IPv6 Prefix Active Bindings: 2
=====
LDP In-Band-SSM IPv6 P2MP Bindings (Active)
=====
Source          Interface
Group           Op          IngLbl  EgrLbl
RootAddr        EgrIf/LspId
EgrNH
-----
No Matching Entries Found
=====
LDP Recursive with In-Band-SSM IPv4 P2MP Bindings (Active)
=====
RootAddr
```

```

InnerRootAddr
Source
Group
EgrNH
Interface
Op
EgrIf/LspId
IngLbl
EgrLbl
-----
No Matching Entries Found
=====
LDP Recursive with In-Band-SSM IPv6 P2MP Bindings (Active)
=====
RootAddr
InnerRootAddr
Source
Group
EgrNH
Interface
Op
EgrIf/LspId
IngLbl
EgrLbl
-----
No Matching Entries Found
=====
LDP In-Band-VPN-SSM IPv4 P2MP Bindings (Active)
=====
Source
Group
RootAddr
EgrNH
RD
Interface
EgrIf/LspId
Op
IngLbl
EgrLbl
-----
No Matching Entries Found
=====
LDP In-Band-VPN-SSM IPv6 P2MP Bindings (Active)
=====
Source
Group
RootAddr
EgrNH
RD
Interface
EgrIf/LspId
Op
IngLbl
EgrLbl
-----
No Matching Entries Found
=====
VPN Recursive with Generic IPv4 P2MP Bindings (Active)
=====
P2MP-Id
RootAddr
InnerRootAddr
EgrNH
Interface
Op
RD
EgrIf/LspId
IngLbl
EgrLbl
-----
No Matching Entries Found
=====
LDP VPN Recursive with Generic IPv6 P2MP Bindings (Active)
=====
P2MP-Id
RootAddr
InnerRootAddr
EgrNH
Interface
Op
RD
EgrIf/LspId
IngLbl
EgrLbl
-----
No Matching Entries Found
=====
LDP GRT Recursive with Generic IPv4 P2MP Bindings (Active)
=====
P2MP-Id
RootAddr
InnerRootAddr
Interface
Op
IngLbl
EgrLbl

```

EgrNH	EgrIf/LspId		

No Matching Entries Found			
=====			
LDP GRT Recursive with Generic IPv6 P2MP Bindings (Active)			
=====			
P2MP-Id	Interface		
RootAddr	Op	IngLbl	EgrLbl
InnerRootAddr			
EgrNH	EgrIf/LspId		

No Matching Entries Found			
=====			
*A: Sar18 Dut-B>show>router>ldp#			

detail

Syntax

detail [**session** *ip-addr* [*label-space*]] [*family*]

Context

show>router>ldp>bindings

Description

This command displays details of LDP bindings.

Parameters

family

displays either IPv4 or IPv6 LDP information

session *ip-addr* [*label-space*]

specifies the IP address and label space identifier

Values *ip-addr* [*label-spa** : *ipv4-address:label-space*
 ipv6-address [*label-space*]
 label-space : 0 to 65535]

Output

The following output is an example of detailed LDP bindings information, and [Table 54: LDP bindings field descriptions](#) describes the fields.

Output example

*A: Sar18 Dut-B>show>router>ldp# bindings active detail	
=====	
LDP Bindings (IPv4 LSR ID 10.20.1.3:0)	
(IPv6 LSR ID 3ffe::a14:103[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
=====
LDP IPv4 Prefix Bindings (Active)
=====
-----
Prefix      : 10.20.1.1/32
Op          : Push
Ing Lbl     : --                Egr Lbl      : 262143
Egr Int/LspId : 1/1/1
EgrNextHop  : 10.10.2.1
Egr. Flags  : None              Ing. Flags : None
Egr If Name : ip-10.10.2.3
Metric      : 1000              Mtu         : 1500
-----
Prefix      : 10.20.1.1/32
Op          : Swap
Ing Lbl     : 262141            Egr Lbl      : 262143
Egr Int/LspId : 1/1/1
EgrNextHop  : 10.10.2.1
Egr. Flags  : None              Ing. Flags : None
Egr If Name : ip-10.10.2.3
Metric      : 1000              Mtu         : 1500
=====
No. of IPv4 Prefix Active Bindings: 2
=====

=====
LDP IPv6 Prefix Bindings (Active)
=====
-----
Prefix      : 3ffe::a14:101/128
Op          : Push
Ing Lbl     : --                Egr Lbl      : 262142
Egr Int/LspId : 1/1/1
EgrNextHop  : fe80::21
Egr. Flags  : None              Ing. Flags : None
Egr If Name : ip-10.10.2.3
Metric      : 1000              Mtu         : 1500
-----
Prefix      : 3ffe::a14:101/128
Op          : Swap
Ing Lbl     : 262136            Egr Lbl      : 262142
Egr Int/LspId : 1/1/1
EgrNextHop  : fe80::21
Egr. Flags  : None              Ing. Flags : None
Egr If Name : ip-10.10.2.3
Metric      : 1000              Mtu         : 1500
-----
No. of IPv6 Prefix Active Bindings: 2
=====

=====
LDP Generic IPv4 P2MP Bindings
=====
No Matching Entries Found
=====

=====
LDP Generic IPv6 P2MP Bindings
=====
No Matching Entries Found
=====

```

```

LDP In-Band-SSM IPv4 P2MP Bindings
=====
No Matching Entries Found
=====

LDP In-Band-SSM IPv6 P2MP Bindings
=====
No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv4 P2MP Bindings
=====
No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv6 P2MP Bindings
=====
No Matching Entries Found
=====

LDP Recursive with In-Band-SSM IPv4 P2MP Bindings
=====
No Matching Entries Found
=====

LDP Recursive with In-Band-SSM IPv6 P2MP Bindings
=====
No Matching Entries Found
=====

LDP VPN Recursive with Generic IPv4 P2MP Bindings
=====
No Matching Entries Found
=====

LDP VPN Recursive with Generic IPv6 P2MP Bindings
=====
No Matching Entries Found
=====

LDP GRT Recursive with Generic IPv4 P2MP Bindings
=====
No Matching Entries Found
=====

LDP GRT Recursive with Generic IPv6 P2MP Bindings
=====
No Matching Entries Found
=====

LDP Service FEC 128 Bindings
=====
Type                : ?-Eth                VcId                : 10
SvcId               : Ukwn                  SdpId               : R. Src
Peer Address        : 1.1.1.1:0
Vc-switching        : No
LMTU                : None                  RMTU                : 1500
Egr. Lbl            : 131070D               Egr. Ctl Word       : No
Egr. Flags          : None                  Egr. Status Bits    : Supported (0x1e)
Egr. Flow Label Tx  : No                    Egr. Flow Label Rx  : No
Egr. PW Status Sig  : Enabled
Egr. Vccv CV Bits   : lsp-ping
Egr. Vccv CC Bits   : router-alert-label
Ing. Lbl            : --                    Ing. Ctl Word        : None

```

```
Ing. Flags      : None          Ing. Status Bits   : N/A
Ing. Flow Label Tx : None       Ing. Flow Label Rx  : None
Ing. Wdraw Reason : None
Ing. PW Status Sig : None
Ing. Vccv CV Bits  : None
Ing. Vccv CC Bits  : None
-----
=====
No. of VC Labels: 1
=====
=====
LDP Service FEC 129 Bindings
=====
No Matching Entries Found
=====
*A: Sar18 Dut-B>show>router>ldp#
```

ipv4

Syntax

ipv4 [**session** *ip-addr*[*label-space*]][**summary** | **detail**]

Context

show>router>ldp>bindings

Description

This command displays LDP active IPv4 bindings.

Parameters

ip-addr[*label-space*]
specifies the IP address and label space identifier

Values *ip-addr*[*label-spa** : *ipv4-address:label-space*
 ipv6-address[*label-space*]
 label-space : 0 to 65535]

detail
displays detailed information

summary
displays information in a summarized format

Output

The following outputs are examples of LDP active IPv4 bindings information, and [Table 54: LDP bindings field descriptions](#) describes the fields.

Output example

```
*A: Sar18 Dut-B>show>router>ldp# bindings ipv4
=====
LDP Bindings (IPv4 LSR ID 10.10.10.10)
```

```

(IPv6 LSR 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
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value), H -
        Hpipe Service
        LF - Lower FEC, UF - Upper FEC
=====
LDP IPv4 Prefix Bindings
=====
Prefix                                IngLbl                                EgrLbl
Peer                                EgrIntf/LspId
EgrNextHop
-----
10.10.10.10/32                        --                                131071
10.10.10.10:0                        1/1/1:100
10.1.1.1
10.12.12.12/32                        131071U                            --
10.10.10.10:0                        --
--
-----
No. of IPv4 Prefix Bindings: 2
=====
LDP Generic IPv4 P2MP Bindings
=====
P2MP-Id
RootAddr                                Interface                                IngLbl    EgrLbl
EgrNH                                EgrIf/LspId
Peer
-----
No Matching Entries Found
=====
LDP In-Band-SSM IPv4 P2MP Bindings
=====
Source
Group
RootAddr                                Interface                                IngLbl    EgrLbl
EgrNH                                EgrIf/LspId
Peer
-----
No Matching Entries Found
=====
LDP In-Band-VPN-SSM IPv4 P2MP Bindings
=====
Source
Group
RootAddr                                RD
EgrNH                                Interface                                IngLbl    EgrLbl
Peer                                EgrIf/LspId
-----
No Matching Entries Found
=====
LDP Recursive with In-Band-SSM IPv4 P2MP Bindings
=====
RootAddr
InnerRootAddr
Source

```

```

Group                               Interface   IngLbl   EgrLbl
EgrNH                               EgrIf/LspId
Peer
-----
No Matching Entries Found
=====
LDP VPN Recursive with Generic IPv4 P2MP Bindings
=====
P2MP-Id                             RD
RootAddr                           Interface
InnerRootAddr                       IngLbl   EgrLbl
EgrNH                               EgrIf/LspId
Peer
-----
No Matching Entries Found
=====
LDP GRT Recursive with Generic IPv4 P2MP Bindings
=====
P2MP-Id                             Interface   IngLbl   EgrLbl
RootAddr                           EgrIf/LspId
InnerRootAddr
EgrNH
Peer
-----
No Matching Entries Found
=====
*A: Sarl8 Dut-B>show>router>ldp#

```

```

*A: Sarl8 Dut-B>show>router>ldp# bindings ipv4 session 10.10.10.10:22 summary
No. of IPv4 Prefix Bindings: 0
No. of Generic IPv4 P2MP Bindings: 0
No. of In-Band-SSM IPv4 P2MP Bindings: 0
No. of In-Band-VPN-SSM IPv4 P2MP Bindings: 0
No. of Recursive with In-Band-SSM IPv4 P2MP Bindings: 0
No. of VPN Recursive with Generic IPv4 P2MP Bindings: 0
No. of GRT Recursive with Generic IPv4 P2MP Bindings: 0
*A: Sarl8 Dut-B>show>router>ldp#

```

```

*A: Sarl8 Dut-B>show>router>ldp# bindings ipv4 detail
=====
LDP Bindings (IPv4 LSR ID 10.10.10.10)
(IPv6 LSR 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
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value), H - Hpipe Service
        LF - Lower FEC, UF - Upper FEC
=====
LDP IPv4 Prefix Bindings
=====
Prefix          : 10.10.10.10/32
-----
Peer            : 10.10.10.10:0
Ing Lbl         : --           Egr Lbl      : 131071
Egr Int/LspId   : 10/10/10:100

```

```

EgrNextHop      : 10.1.1.1
Egr. Flags      : None                Ing. Flags : None
Egr If Name     : toA
Metric          : 1                   Mtu         : 1554
-----
Prefix          : 10.12.12.12/32
-----
Peer            : 10.10.10.1:0
Ing Lbl         : 131071U             Egr Lbl      : --
Egr Int/LspId   : --
EgrNextHop      : --
Egr. Flags      : None                Ing. Flags : None
=====
No. of IPv4 Prefix Bindings: 2
=====
LDP Generic IPv4 P2MP Bindings
=====
No Matching Entries Found
=====
LDP In-Band-SSM IPv4 P2MP Bindings
=====
No Matching Entries Found
=====
LDP In-Band-VPN-SSM IPv4 P2MP Bindings
=====
No Matching Entries Found
=====
LDP Recursive with In-Band-SSM IPv4 P2MP Bindings
=====
No Matching Entries Found
=====
LDP VPN Recursive with Generic IPv4 P2MP Bindings
=====
No Matching Entries Found
=====
LDP GRT Recursive with Generic IPv4 P2MP Bindings
=====
No Matching Entries Found
=====
*A: Sar18 Dut-B>show>router>ldp#

```

ipv6

Syntax

ipv6 [**session** *ip-addr* [*label-space*]] [**summary** | **detail**]

Context

show>router>ldp>bindings

Description

This command displays LDP active IPv6 bindings.

Parameters

ip-addr[*label-space*]
specifies the IP address and label space identifier

Values *ip-addr*[*label-spa** : *ipv4-address*:*label-space*
 ipv6-address[*label-space*]
 label-space : 0 to 65535]

detail
displays detailed information

summary
displays information in a summarized format

Output

The following output is an example of LDP active IPv6 bindings information, and [Table 54: LDP bindings field descriptions](#) describes the fields.

Output example

```
*A:Sar18 Dut-B# show router ldp bindings ipv6
=====
LDP Bindings (IPv4 LSR ID 10.12.12.12)
(IPv6 LSR 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
A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value), H -
Hpipe Service
LF - Lower FEC, UF - Upper FEC
=====
LDP IPv6 Prefix Bindings
=====
Prefix                               IngLbl                               EgrLbl
Peer                                EgrIntf/LspId
EgrNextHop
-----
No Matching Entries Found
=====
LDP Generic IPv6 P2MP Bindings
=====
P2MP-Id
RootAddr                             Interface                             IngLbl    EgrLbl
EgrNH                                EgrIf/LspId
Peer
-----
No Matching Entries Found
=====
LDP In-Band-SSM IPv6 P2MP Bindings
=====
Source
Group                                Interface                             IngLbl    EgrLbl
RootAddr
```

EgrNH	EgrIf/LspId		
Peer			

No Matching Entries Found			
=====			
LDP In-Band-VPN-SSM IPv6 P2MP Bindings			
=====			
Source	RD		
Group	Interface	IngLbl	EgrLbl
RootAddr	EgrIf/LspId		
EgrNH			
Peer			

No Matching Entries Found			
=====			
LDP Recursive with In-Band-SSM IPv6 P2MP Bindings			
=====			
RootAddr			
InnerRootAddr			
Source	Interface	IngLbl	EgrLbl
Group	EgrIf/LspId		
EgrNH			
Peer			

No Matching Entries Found			
=====			
LDP VPN Recursive with Generic IPv6 P2MP Bindings			
=====			
P2MP-Id	RD		
RootAddr	Interface		
InnerRootAddr		IngLbl	EgrLbl
EgrNH	EgrIf/LspId		
Peer			

No Matching Entries Found			
=====			
LDP GRT Recursive with Generic IPv6 P2MP Bindings			
=====			
P2MP-Id			
RootAddr	Interface	IngLbl	EgrLbl
InnerRootAddr			
EgrNH	EgrIf/LspId		
Peer			

No Matching Entries Found			
=====			
*A: Sar18 Dut-B#			

label-type

Syntax

label-type start-label *start-label* [end-label *end-label*] label-type [*family*]

Context

show>router>ldp>bindings

Description

This command displays LDP FEC bindings by matching labels.

Parameters

- start-label*

specifies a label value to begin the display

Values16 to 1048575
- end-label*

specifies a label value to end the display

Values17 to 1048575
- label-type*

specifies a label type to display

Valuesingress-label or egress-label
- family*

displays either IPv4 or IPv6 LDP information

Valuesipv4 or ipv6

Output

The following output is an example of LDP FEC bindings information by matching labels, and [Table 54: LDP bindings field descriptions](#) describes the fields.

Output example

```
*A:Sar18 Dut-B>show>router>ldp# bindings label-type start-label 16 end-label 99 ingress-label
=====
LDP Bindings (IPv4 LSR ID 10.12.12.12)
              (IPv6 LSR ID ::)
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        LF - Lower FEC, UF - Upper FEC
=====
LDP IPv4 Prefix Bindings
=====
Prefix                               IngLbl                               EgrLbl
Peer                                EgrIntf/LspId
EgrNextHop
-----
No Matching Entries Found
=====
LDP IPv6 Prefix Bindings
=====
Prefix                               IngLbl                               EgrLbl
Peer                                EgrIntf/LspId
EgrNextHop
-----
No Matching Entries Found
```

```
=====
LDP Service FEC 128 Bindings
=====
Type                               VCId      SDPIId    IngLbl   LMTU
Peer                               SvcId     EgrLbl    RMTU
-----
No Matching Entries Found
=====

LDP Service FEC 129 Bindings
=====
SAII                               AGII       IngLbl    LMTU
TAII                               Type       EgrLbl    RMTU
Peer                               SvcId     SDPIId
-----
No Matching Entries Found
=====
*A: Sar18 Dut-B>show>router>ldp#
```

p2mp

Syntax

```
p2mp p2mp-id identifier root ip-address [session ip-addr [label-space]] [summary | detail]
p2mp [session ip-addr [label-space]] [family] [summary | detail] [opaque-type opaque-type]
p2mp source ip-address group mcast-address root ip-address [session ip-addr [label-space]][family]
[summary | detail] inner-root ip-address
p2mp source ip-address group mcast-address root ip-address [rd rd] [session ip-addr [label-space]]
[summary | detail]
p2mp source ip-address group mcast-address [session ip-addr [label-space]] [family] [summary | detail]
[innermost-root ip-address]
```

Context

```
show>router>ldp>bindings
```

Description

This command displays LDP P2MP FEC bindings.

Parameters

- detail**
displays detailed information
- summary**
displays information in a summarized format
- family*
displays either IPv4 or IPv6 active LDP information
- group mcast-address**
displays the P2MP group multicast address bindings

- inner-root** *ip-address*
displays recursive FECs whose inner root address matches the specified address
- innermost-root** *ip-address*
displays recursive FECs whose inner root address matches the specified address and non-recursive FECs that have a root address that matches the specified address
- opaque-type** *opaque-type*
specifies the type of a multipoint opaque value element

Values generic, ssm, vpn-ssm, recursive-ssm, vpn-recursive, grt-recursive
- p2mp-id** *identifier*
displays LDP active P2MP identifier bindings

Values 0 to 4294967295
- rd** *rd*
displays information for the route distinguisher

Values *ip-addr:comm-val | 2byte-asnumber:ext-comm-val | 4byte-asnumber:comm-val*
- root** *ip-address*
displays root IP address information
- session** *ip-addr[label-space]*
displays information for the LDP session IP address and label space

Values *ipv4-address:label-space*
 ipv6-address[label-space]
 label-space: 0 to 65535
- source** *ip-address*
displays LDP active P2MP source bindings

Output

The following output is an example of LDP P2MP FEC bindings information, and [Table 54: LDP bindings field descriptions](#) describes the fields.

Output example

```
*A:7705:Dut-F# show router ldp bindings p2mp detail
=====
LDP Bindings (IPv4 LSR ID 10.20.1.6)
(IPv6 LSR ID ::)
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        LF - Lower FEC, UF - Upper FEC
=====
LDP Generic IPv4 P2MP Bindings
=====
-----
P2MP Type      : 1                      P2MP-Id      : 8193
```

```

Root-Addr      : 10.20.1.1
-----
Peer           : 10.20.1.4:0
Ing Lbl        : 131062U
Egr Lbl        : --
Egr Int/LspId  : --
EgrNextHop     : --
Egr. Flags     : None           Ing. Flags : None
-----
P2MP Type      : 1              P2MP-Id   : 8193
Root-Addr      : 10.20.1.1
-----
Peer           : 10.20.1.5:0
Ing Lbl        : --
Egr Lbl        : 131062
Egr Int/LspId  : 1/1/3:0
EgrNextHop     : 10.180.10.5
Egr. Flags     : None           Ing. Flags : None
Egr If Name    : ip-10.180.10.6
Metric         : 1              Mtu        : 1496
-----
P2MP Type      : 1              P2MP-Id   : 8193
Root-Addr      : 10.20.1.2
-----
Peer           : 10.20.1.4:0
Ing Lbl        : 131064U
Egr Lbl        : --
Egr Int/LspId  : --
EgrNextHop     : --
Egr. Flags     : None           Ing. Flags : None
-----
P2MP Type      : 1              P2MP-Id   : 8193
Root-Addr      : 10.20.1.2
-----
Peer           : 10.20.1.5:0
Ing Lbl        : --
Egr Lbl        : 131064
Egr Int/LspId  : 1/1/3:0
EgrNextHop     : 10.180.10.5
Egr. Flags     : None           Ing. Flags : None
Egr If Name    : ip-10.180.10.6
Metric         : 1              Mtu        : 1496
-----
P2MP Type      : 1              P2MP-Id   : 8193
Root-Addr      : 10.20.1.3
-----
Peer           : 10.20.1.4:0
Ing Lbl        : 131063U
Egr Lbl        : --
Egr Int/LspId  : --
EgrNextHop     : --
Egr. Flags     : None           Ing. Flags : None
-----
P2MP Type      : 1              P2MP-Id   : 8193
Root-Addr      : 10.20.1.3
-----
Peer           : 10.20.1.5:0
Ing Lbl        : --
Egr Lbl        : 131063
Egr Int/LspId  : 1/1/3:0
EgrNextHop     : 10.180.10.5
Egr. Flags     : None           Ing. Flags : None
Egr If Name    : ip-10.180.10.6
Metric         : 1              Mtu        : 1496

```

```

-----
P2MP Type      : 1                P2MP-Id      : 8193
Root-Addr      : 10.20.1.5
-----
Peer           : 10.20.1.4:0
Ing Lbl        : --
Egr Lbl        : 131065
Egr Int/LspId  : 1/1/1:0
EgrNextHop     : 10.180.9.4
Egr. Flags     : None              Ing. Flags : None
Egr If Name    : ip-10.180.9.6
Metric         : 1                Mtu          : 1496
-----
P2MP Type      : 1                P2MP-Id      : 8193
Root-Addr      : 10.20.1.5
-----
Peer           : 10.20.1.5:0
Ing Lbl        : 131065U
Egr Lbl        : --
Egr Int/LspId  : --
EgrNextHop     : --
Egr. Flags     : None              Ing. Flags : None
-----
P2MP Type      : 1                P2MP-Id      : 8193
Root-Addr      : 10.20.1.6
-----
Peer           : 10.20.1.4:0
Ing Lbl        : --
Egr Lbl        : 131061
Egr Int/LspId  : 1/1/1:0
EgrNextHop     : 10.180.9.4
Egr. Flags     : None              Ing. Flags : None
Egr If Name    : ip-10.180.9.6
Metric         : 1                Mtu          : 1496
-----
P2MP Type      : 1                P2MP-Id      : 8193
Root-Addr      : 10.20.1.6
-----
Peer           : 10.20.1.5:0
Ing Lbl        : --
Egr Lbl        : 131061
Egr Int/LspId  : 1/1/3:0
EgrNextHop     : 10.180.10.5
Egr. Flags     : None              Ing. Flags : None
Egr If Name    : ip-10.180.10.6
Metric         : 1                Mtu          : 1496
=====
No. of Generic IPv4 P2MP Bindings: 10
=====
LDP Generic IPv6 P2MP Bindings
=====
No Matching Entries Found
=====
LDP In-Band-SSM IPv4 P2MP Bindings
=====
No Matching Entries Found
=====
LDP In-Band-SSM IPv6 P2MP Bindings
=====
No Matching Entries Found
=====

```

```

=====
LDP In-Band-VPN-SSM IPv4 P2MP Bindings
=====
No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv6 P2MP Bindings
=====
No Matching Entries Found
=====

LDP Recursive with In-Band-SSM IPv4 P2MP Bindings
=====
No Matching Entries Found
=====

LDP Recursive with In-Band-SSM IPv6 P2MP Bindings
=====
No Matching Entries Found
=====

LDP VPN Recursive with Generic IPv4 P2MP Bindings
=====
No Matching Entries Found
=====

LDP VPN Recursive with Generic IPv6 P2MP Bindings
=====
No Matching Entries Found
=====

LDP GRT Recursive with Generic IPv4 P2MP Bindings
=====
No Matching Entries Found
=====

LDP GRT Recursive with Generic IPv6 P2MP Bindings
=====
No Matching Entries Found
=====
*A:7705:Dut-F#

```

prefixes

Syntax

prefixes prefix *ip-prefix/ip-prefix-length* [**summary** | **detail**] [**session** *ip-addr[label-space]*]

prefixes [*family*] [**summary** | **detail**] [**session** *ip-addr[label-space]*]

Context

show>router>ldp>bindings

Description

This command displays LDP prefix FEC bindings.

Parameters

- prefix *ip-prefix**ip-prefix-length*

displays information for the specified IP prefix and mask length
- detail

displays detailed information
- summary

displays information in a summarized format
- family

displays either IPv4 or IPv6 active LDP information

Values

ipv4 or ipv6
- session *ip-addr*

displays configuration information about LDP sessions
- label-space

specifies the label space identifier that the router is advertising on the interface

Values

0 to 65535

Output

The following output is an example of LDP prefix FEC bindings information, and [Table 54: LDP bindings field descriptions](#) describes the fields.

Output example

```
*A:Sar18 Dut-B>show>router>ldp# bindings prefixes
=====
LDP Bindings (IPv4 LSR ID 10.12.12.2)
(IPv6 LSR ID ::)
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
       LF - Lower FEC, UF - Upper FEC
=====
LDP IPv4 Prefix Bindings
=====
Prefix                               IngLbl                               EgrLbl
Peer                               EgrIntf/LspId
EgrNextHop
-----
10.10.10.1/32                        --                                131071
1.1.1.1:0                          1/1/1:100
10.1.1.1
10.12.12.12/32                     131071U
10.1.1.1:0                         --
--
-----
No. of IPv4 Prefix Bindings: 2
=====
LDP IPv6 Prefix Bindings
=====
Prefix                               IngLbl                               EgrLbl
Peer                               EgrIntf/LspId
```

```

EgrNextHop
-----
No Matching Entries Found
=====
*A: Sar18 Dut-B>show>router>ldp#

*A: Sar18 Dut-B>show>router>ldp# bindings prefixes detail
=====
LDP Bindings (IPv4 LSR ID 10.12.12.2)
(IPv6 LSR ID ::)
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
        LF - Lower FEC, UF - Upper FEC
=====
LDP IPv4 Prefix Bindings
=====
Prefix          : 10.10.10.1/32
-----
Peer            : 10.10.10.1:0
Ing Lbl         : --                Egr Lbl      : 131071
Egr Int/LspId   : 10/10/10:100
EgrNextHop      : 10.1.1.1
Egr. Flags      : None              Ing. Flags  : None
Egr If Name     : toA
Metric          : 1                 Mtu         : 1554
-----
Prefix          : 10.12.12.2/32
-----
Peer            : 10.10.10.1:0
Ing Lbl         : 131071U           Egr Lbl      : --
Egr Int/LspId   : --
EgrNextHop      : --
Egr. Flags      : None              Ing. Flags  : None
=====
No. of IPv4 Prefix Bindings: 2
=====
LDP IPv6 Prefix Bindings
=====
No Matching Entries Found
=====
*A: Sar18 Dut-B>show>router>ldp#

```

services

Syntax

services vc-type *vc-type* **saii** *global-id:prefix:ac-id taii* [256 chars max] **agi** *agi* [**detail**] [**service-id** *service-id*] [**session** *ip-addr[label-space]*]

services vc-type *vc-type* **agi** *agi* [**detail**] [**service-id** *service-id*] [**session** *ip-addr[label-space]*]

services [**vc-type** *vc-type*] [**svc-fec-type**] [**detail**] [**service-id** *service-id*] [**session** *ip-addr[label-space]*]

services vc-type *vc-type* **vc-id** *vc-id* [**detail**] [**service-id** *service-id*] [**session** *ip-addr[label-space]*]

Context

show>router>ldp>bindings

Description

This command displays LDP service FEC bindings.

Parameters**vc-type** *vc-type*

displays information for the VC type associated with this service FEC

Values ethernet, vlan, mirror, frdlci, atmsdu, atmcell, atmvc, atmvp, ipipe, satop-e1, satop-t1, cesopsn, cesopsn-cas

vc-id *vc-id*

displays information for the VC ID associated with this service FEC

Values 1 to 4294967295

saii *global-id:prefix:ac-id*

specifies the SAI (source attachment individual identifier)

Values *number:number* | *a.b.c.d:number*

taii

specifies the TAI (target attachment individual identifier), up to 256 characters, associated with this service FEC

svc-fec-type

specifies the FEC type

Values fec128, fec129

agi *agi*

specifies the attachment group identifier TLV associated with this service FEC

Values *ip-addr:comm-val* | *2byte-asnumber:ext-comm-val* | *4byte-asnumber:comm-val*

ip-addr : a.b.c.d

comm-val : 0 to 65535

2byte-asnumber : 1 to 65535

ext-comm-val : 0 to 4294967295

4byte-asnumber : 1 to 4294967295

null - means all value is 0

detail

displays detailed information

service-id

specifies the service ID number to display

Values 1 to 2148007980 | *svc-name* (64 char max)

svc-fec-type
specifies the FEC type

Values fec128, fec129

ip-addr[label-space]
specifies the IP address and the label space identifier that the router is advertising on the interface

Values *ipv4-address:label-space*
ipv6-address[label-space]
label-space : 0 to 65535

Output

The following outputs are examples of LDP service FEC bindings information, and [Table 54: LDP bindings field descriptions](#) describes the fields.

Output example

```
*A:Sar18 Dut-B>show>router>ldp# bindings services vc-type ethernet vc-id 999
=====
LDP Bindings (IPv4 LSR ID 10.2.2.2)
(IPv6 LSR 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
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value), H -
        Hpipe Service
        LF - Lower FEC, UF - Upper FEC
=====
LDP Service FEC 128 Bindings
=====
Type          VCId      SDPId      IngLbl  LMTU
Peer          SvcId      EgrLbl  RMTU
-----
No Matching Entries Found
=====

*A:Sar18 Dut-B>show>router>ldp# bindings services vc-type ethernet vc-
id 999 service-id 7777
=====
LDP Bindings (IPv4 LSR ID 10.2.2.2)
(IPv6 LSR 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
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value), H -
        Hpipe Service
        LF - Lower FEC, UF - Upper FEC
```

```
=====
LDP Service FEC 128 Bindings
=====
Type                               VCId      SDPIId    IngLbl   LMTU
Peer                               SvcId      EgrLbl    RMTU
-----
No Matching Entries Found
=====

*A:Sar18 Dut-B>show>router>ldp# bindings services vc-type ethernet vc-
id 999 service-id 7777 session 10.1.1.1:333
=====
LDP Bindings (IPv4 LSR ID 10.2.2.2)
(IPv6 LSR 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
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value), H -
        Hpipe Service
        LF - Lower FEC, UF - Upper FEC
=====
LDP Service FEC 128 Bindings
=====
Type                               VCId      SDPIId    IngLbl   LMTU
Peer                               SvcId      EgrLbl    RMTU
-----
No Matching Entries Found
=====
*A:Sar18 Dut-B>show>router>ldp#
```

session

Syntax
`session [family] [summary | detail] ip-addr[label-space]`

Context
show>router>ldp>bindings

Description
This command displays LDP FEC bindings by matching peer LSR ID.

- Parameters**
- detail**
displays detailed information
 - summary**
displays information in a summarized format
 - family**
displays either IPv4 or IPv6 LDP session information

Values ipv4 or ipv6

ip-addr[label-space]

specifies the IP address and the label space identifier that the router is advertising on the interface

Values *ipv4-address:label-space*

ipv6-address[label-space]

label-space : 0 to 65535

Output

The following outputs are examples of LDP P2MP FEC bindings information by matching peer LSR ID, and [Table 54: LDP bindings field descriptions](#) describes the fields.

Output example

```
*A:Dut-A# show router ldp bindings session 3ffe::a14:103 summary
```

```
No. of IPv4 Prefix Bindings: 0
No. of IPv6 Prefix Bindings: 6
No. of Generic IPv4 P2MP Bindings: 0
No. of Generic IPv6 P2MP Bindings: 0
No. of In-Band-SSM IPv4 P2MP Bindings: 0
No. of In-Band-SSM IPv6 P2MP Bindings: 0
No. of In-Band-VPN-SSM IPv4 P2MP Bindings: 0
No. of In-Band-VPN-SSM IPv6 P2MP Bindings: 0
No. of VC Labels: 0
No. of FEC 129s: 0
```

```
*A:Dut-A# show router ldp bindings session 3ffe::a14:103 detail
```

```
=====
LDP Bindings (IPv4 LSR ID 10.20.1.1:0)
              (IPv6 LSR ID 3ffe::a14:101[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
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
        BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value)
=====
```

```
LDP IPv4 Prefix Bindings
=====
```

```
No Matching Entries Found
=====
```

```
LDP IPv6 Prefix Bindings
=====
```

```
Prefix       : 3ffe::a14:101/128
-----
```

```
Peer         : 3ffe::a14:103[0]
Ing Lbl      : 262142U           Egr Lbl      :  --
Egr Int/LspId :  --
EgrNextHop   :  --
Egr. Flags   : None             Ing. Flags : None
-----
```

```
Prefix       : 3ffe::a14:102/128
-----
```

```

Peer      : 3ffe::a14:103[0]
Ing Lbl   : 262136U          Egr Lbl   : 262138
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None           Ing. Flags : None
Egr If Name : n/a

```

```

-----
Prefix    : 3ffe::a14:103/128
-----

```

```

Peer      : 3ffe::a14:103[0]
Ing Lbl   : --              Egr Lbl   : 262142
Egr Int/LspId : 1/1/2
EgrNextHop : fe80::23
Egr. Flags : None           Ing. Flags : None
Egr If Name : ip-10.10.2.1
Metric    : 1000            Mtu       : 1500

```

```

-----
Prefix    : 3ffe::a14:104/128
-----

```

```

Peer      : 3ffe::a14:103[0]
Ing Lbl   : 262132U          Egr Lbl   : 262134
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None           Ing. Flags : None
Egr If Name : n/a

```

```

-----
Prefix    : 3ffe::a14:105/128
-----

```

```

Peer      : 3ffe::a14:103[0]
Ing Lbl   : 262134N          Egr Lbl   : 262132
Egr Int/LspId : 1/1/2
EgrNextHop : fe80::23
Egr. Flags : None           Ing. Flags : None
Egr If Name : ip-10.10.2.1
Metric    : 2000            Mtu       : 1500

```

```

-----
Prefix    : 3ffe::a14:106/128
-----

```

```

Peer      : 3ffe::a14:103[0]
Ing Lbl   : 262133U          Egr Lbl   : 262133
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None           Ing. Flags : None
Egr If Name : n/a

```

```

=====
No. of IPv6 Prefix Bindings: 6
=====

```

```

=====
LDP Generic IPv4 P2MP Bindings
=====

```

```

No Matching Entries Found
=====

```

```

=====
LDP Generic IPv6 P2MP Bindings
=====

```

```

No Matching Entries Found
=====

```

```

=====
LDP In-Band-SSM IPv4 P2MP Bindings
=====

```

```

No Matching Entries Found
=====

```

```

=====
LDP In-Band-SSM IPv6 P2MP Bindings
=====

```

```
=====
No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv4 P2MP Bindings
=====
No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv6 P2MP Bindings
=====
No Matching Entries Found
=====

LDP Service FEC 128 Bindings
=====
No Matching Entries Found
=====

LDP Service FEC 129 Bindings
=====
No Matching Entries Found
=====

*A:Dut-A#

*A:Dut-A# show router ldp bindings session 10.20.1.3 ipv4
=====
LDP Bindings (IPv4 LSR ID 10.20.1.1:0)
              (IPv6 LSR ID 3ffe::a14:101[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
       A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
       P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
       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                         262143U                               --
10.20.1.3:0                          --
--

10.20.1.2/32                         262141U                               262140
10.20.1.3:0                          --
--

10.20.1.3/32                         --                                   262143
10.20.1.3:0                          1/1/2
10.10.2.3

10.20.1.4/32                         262139U                               262139
10.20.1.3:0                          --
--

10.20.1.5/32                         262138N                               262137
10.20.1.3:0                          1/1/2
10.10.2.3
```

```
10.20.1.6/32                262135U                262135
10.20.1.3:0                --
--

-----
No. of IPv4 Prefix Bindings: 6
=====
LDP Generic IPv4 P2MP Bindings
=====
P2MP-Id
RootAddr                    Interface      IngLbl    EgrLbl
EgrNH                      EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

=====
LDP In-Band-SSM IPv4 P2MP Bindings
=====
Source
Group
RootAddr                    Interface      IngLbl    EgrLbl
EgrNH                      EgrIf/LspId
Peer
-----
No Matching Entries Found
=====

=====
LDP In-Band-VPN-SSM IPv4 P2MP Bindings
=====
Source
Group                        RD
RootAddr                    Interface      IngLbl    EgrLbl
EgrNH                      EgrIf/LspId
Peer
-----
No Matching Entries Found
=====
*A:Dut-A#
```

summary

Syntax
summary [session *ip-addr*[*label-space*]] [**ipv4** | **ipv6**]

Context
show>router>ldp>bindings

Description
This command displays a summary of LDP bindings.

Parameters

session *ip-addr[label-space]*

specifies the IP address and label space identifier

Values *ip-addr[label-spa* : ipv4-address:label-space ipv6-address[label-space]*
label-space: 0 to 65535

ipv4

displays IPv4 summary bindings information

ipv6

displays IPv6 summary bindings information

Output

The following output is an example of summary LDP bindings information, and [Table 54: LDP bindings field descriptions](#) describes the fields.

Output example

```
*A:Sar18 Dut-B>show>router>ldp# bindings summary
No. of IPv4 Prefix Bindings: 2
No. of IPv6 Prefix Bindings: 0
No. of Generic IPv4 P2MP Bindings: 0
No. of Generic IPv6 P2MP Bindings: 0
No. of In-Band-SSM IPv4 P2MP Bindings: 0
No. of In-Band-SSM IPv6 P2MP Bindings: 0
No. of In-Band-VPN-SSM IPv4 P2MP Bindings: 0
No. of In-Band-VPN-SSM IPv6 P2MP Bindings: 0
No. of Recursive with In-Band-SSM IPv4 P2MP Bindings: 0
No. of Recursive with In-Band-SSM IPv6 P2MP Bindings: 0
No. of VPN Recursive with Generic IPv4 P2MP Bindings: 0
No. of VPN Recursive with Generic IPv6 P2MP Binding: 0
No. of GRT Recursive with Generic IPv4 P2MP Bindings: 0
No. of GRT Recursive with Generic IPv6 P2MP Binding: 0
No. of VC Labels: 1
No. of FEC 129s: 0
*A:Sar18 Dut-B>show>router>ldp# bindings summary
```

6.12.2.3 Clear commands

fec-egress-statistics

Syntax

fec-egress-statistics [*ip-prefix/mask*]

Context

clear>router>ldp

Description

This command clears LDP FEC statistics.



Note: When LDP FEC statistics are cleared, the current aggregate statistics count is recorded as a baseline and is used to provide a relative count each time the statistics are viewed with the **show** command. Because this baseline number is not reconciled between the active and inactive CSMs, after a CSM activity switch the statistics on the newly active CSM shows the aggregate count as though no clear command has been executed.

Parameters

ip-prefix/mask

the IP prefix and prefix length associated with the prefix FEC

instance

Syntax

instance

Context

clear>router>ldp

Description

This command resets the LDP instance.

interface

Syntax

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

Context

clear>router>ldp

Description

This command restarts or clears statistics for LDP interfaces.

Parameters

ip-int-name

specifies an existing interface. If the string contains special characters (such as #, \$, spaces), the entire string must be enclosed within double quotes.

statistics

clears only the statistics for an interface

peer

Syntax

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

Context

clear>router>ldp

Description

This command restarts or clears statistics for LDP targeted peers.

Parameters

ip-address

specifies a targeted peer

statistics

clears only the statistics for a targeted peer

session

Syntax

session *ip-addr* [:*label-space*][**statistics**]

Context

clear>router>ldp

Description

This command restarts or clears statistics for LDP sessions.

Parameters

ip-addr

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

statistics

clears only the statistics for a session

statistics

Syntax
statistics

Context
clear>router>ldp

Description
This command clears LDP instance statistics.

6.12.2.4 Monitor commands

fec-egress-stats

Syntax
fec-egress-stats *ip-prefix/mask* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]

Context
monitor>router>ldp

Description
This command displays egress statistics for LDP FEC prefixes at the configured interval until the configured count is reached.

Monitor commands are similar to **show** commands, but only statistical information is displayed. Monitor commands display the selected statistics according to the configured number of times at the interval specified.

Parameters
ip-prefix/mask
the IP prefix and prefix length associated with the prefix FEC
seconds
specifies the interval for each display, in seconds
Values 3 to 60
Default 10
repeat
specifies how many times the command is repeated
Values 1 to 999

Default 10

absolute

displays raw statistics, without processing. No calculations are performed on the delta or rate statistics.

rate

displays the rate per second for each statistic instead of the delta

Output

The following output is an example of statistics for an LDP FEC instance.

Output example

```
=====
Monitor egress statistics for LDP FEC prefix 10.10.10.29/32
=====
-----
At time t = 0 sec (Base Statistics)
-----
FEC Prefix/Mask      : 10.10.10.29/32
-----
Collect Stats       : Disabled           Accounting Plcy.    : None
Admin State        : Up
FC BE
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets      : 0                   OutProf Octets     : 0
FC L2
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets      : 0                   OutProf Octets     : 0
FC AF
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets      : 0                   OutProf Octets     : 0
FC L1
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets      : 0                   OutProf Octets     : 0
FC H2
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets      : 0                   OutProf Octets     : 0
FC EF
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets      : 0                   OutProf Octets     : 0
FC H1
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets      : 0                   OutProf Octets     : 0
FC NC
InProf Pkts        : 0                   OutProf Pkts       : 0
InProf Octets      : 0                   OutProf Octets     : 0
-----
At time t = 3 sec (Mode: Absolute)
-----
=====
Monitor egress statistics for LDP FEC prefix 10.10.10.29/32
=====
-----
At time t = 0 sec (Base Statistics)
-----
FEC Prefix/Mask      : 10.10.10.29/32
-----
```

Collect Stats	: Disabled	Accounting Plcy.	: None
Admin State	: Up		
FC BE			
InProf Pkts	: 0	OutProf Pkts	: 0
InProf Octets	: 0	OutProf Octets	: 0
FC L2			
InProf Pkts	: 0	OutProf Pkts	: 0
InProf Octets	: 0	OutProf Octets	: 0
FC AF			
InProf Pkts	: 0	OutProf Pkts	: 0
InProf Octets	: 0	OutProf Octets	: 0
FC L1			
InProf Pkts	: 0	OutProf Pkts	: 0
InProf Octets	: 0	OutProf Octets	: 0
FC H2			
InProf Pkts	: 0	OutProf Pkts	: 0
InProf Octets	: 0	OutProf Octets	: 0
FC EF			
InProf Pkts	: 0	OutProf Pkts	: 0
InProf Octets	: 0	OutProf Octets	: 0
FC H1			
InProf Pkts	: 0	OutProf Pkts	: 0
InProf Octets	: 0	OutProf Octets	: 0
FC NC			
InProf Pkts	: 0	OutProf Pkts	: 0
InProf Octets	: 0	OutProf Octets	: 0

At time t = 3 sec (Mode: Rate)			

session

Syntax

session*/ldp-id*[*/ldp-id...*(up to 5 max)][*intervalseconds*][*repeatrepeat*][*absolute* | *rate*]

Context

monitor>router>ldp

Description

This command displays statistical information for LDP sessions at the configured interval until the configured count is reached.

The first screen displays the current statistics related to the specified LDP sessions. The subsequent statistical information listed for each interval is displayed as a delta to the previous screen output. When the keyword **rate** is specified, the rate per second for each statistic is displayed instead of the delta.

Monitor commands are similar to **show** commands, but only statistical information is displayed. Monitor commands display the selected statistics according to the configured number of times at the interval specified.

Parameters

ldp-id
specifies the IP address of the LDP session to display

Values *ip-addr[:label-space]*
 ip-addr – a.b.c.d
 label-space – [0 to 65535]

seconds
 configures the interval for each display in seconds

Values 3 to 60

Default 10

repeat
 configures how many times the command is repeated

Values 1 to 999

Default 10

absolute
 displays raw statistics, without processing. No calculations are performed on the delta or rate statistics.

rate
 displays the rate per second for each statistic instead of the delta

Output
The following output is an example of statistical information for the LDP session.

Output example

```
ALU-103>monitor>router>ldp# session 10.10.10.104 interval 3 repeat 3 absolute
=====
Monitor statistics for LDP Session 10.10.10.104
=====

```

	Sent	Received

At time t = 0 sec (Base Statistics)		

FECs	1	2
Hello	5288	5289
Keepalive	8225	8225
Init	1	1
Label Mapping	1	4
Label Request	0	0
Label Release	0	0
Label Withdraw	0	0
Label Abort	0	0
Notification	0	0
Address	1	1
Address Withdraw	0	0

At time t = 3 sec (Mode: Absolute)		

FECs	1	2
Hello	5288	5289
Keepalive	8226	8226

```

Init          1          1
Label Mapping 1          4
Label Request 0          0
Label Release 0          0
Label Withdraw 0         0
Label Abort   0          0
Notification  0          0
Address       1          1
Address Withdraw 0        0
-----
At time t = 6 sec (Mode: Absolute)
-----
FECs          1          2
Hello         5288       5290
Keepalive     8226       8226
Init          1          1
Label Mapping 1          4
Label Request 0          0
Label Release 0          0
Label Withdraw 0         0
Label Abort   0          0
Notification  0          0
Address       1          1
Address Withdraw 0        0
-----
At time t = 9 sec (Mode: Absolute)
-----
FECs          1          2
Hello         5288       5290
Keepalive     8226       8226
Init          1          1
Label Mapping 1          4
Label Request 0          0
Label Release 0          0
Label Withdraw 0         0
Label Abort   0          0
Notification  0          0
Address       1          1
Address Withdraw 0        0
=====
ALU-12>monitor>router>ldp#

```

```

ALU-12>monitor>router>ldp# session 10.10.10.104 interval 3 repeat 3 rate
=====
Monitor statistics for LDP Session 10.10.10.104
=====
Sent          Received
-----
At time t = 0 sec (Base Statistics)
-----
FECs          1          2
Hello         5289       5290
Keepalive     8227       8227
Init          1          1
Label Mapping 1          4
Label Request 0          0
Label Release 0          0
Label Withdraw 0         0
Label Abort   0          0
Notification  0          0
Address       1          1
Address Withdraw 0        0
-----

```

At time t = 3 sec (Mode: Rate)		

FECs	0	0
Hello	0	0
Keepalive	0	0
Init	0	0
Label Mapping	0	0
Label Request	0	0
Label Release	0	0
Label Withdraw	0	0
Label Abort	0	0
Notification	0	0
Address	0	0
Address Withdraw	0	0

At time t = 6 sec (Mode: Rate)		

FECs	0	0
Hello	0	0
Keepalive	0	0
Init	0	0
Label Mapping	0	0
Label Request	0	0
Label Release	0	0
Label Withdraw	0	0
Label Abort	0	0
Notification	0	0
Address	0	0
Address Withdraw	0	0

At time t = 9 sec (Mode: Rate)		

FECs	0	0
Hello	0	0
Keepalive	0	0
Init	0	0
Label Mapping	0	0
Label Request	0	0
Label Release	0	0
Label Withdraw	0	0
Label Abort	0	0
Notification	0	0
Address	0	0
Address Withdraw	0	0
=====		
ALU-12>monitor>router>ldp#		

statistics

Syntax
statistics [*interval seconds*][*repeat repeat*][*absolute* | *rate*]

Context
monitor>router>ldp

Description

This command displays statistics for an LDP instance at the configured interval until the configured count is reached.

The first screen displays the current statistics related to the LDP statistics. The subsequent statistical information listed for each interval is displayed as a delta to the previous screen output.

When the keyword **rate** is specified, the rate per second for each statistic is displayed instead of the delta.

Monitor commands are similar to **show** commands, but only statistical information is displayed. Monitor commands display the selected statistics according to the configured number of times at the interval specified.

Parameters

seconds

configures the interval for each display in seconds

Values 3 to 60

Default 10

repeat

configures how many times the command is repeated

Values 1 to 999

Default 10

absolute

displays raw statistics, without processing. No calculations are performed on the delta or rate statistics.

rate

displays the rate per second for each statistic instead of the delta

Output

The following outputs are examples of statistics for an LDP instance.

Output example

```
ALU-12>monitor>router>ldp# statistics interval 3 repeat 3 absolute
=====
Monitor statistics for LDP instance
=====
At time t = 0 sec (Base Statistics)
-----
Addr FECs Sent      : 0                      Addr FECs Recv      : 0
Serv FECs Sent      : 1                      Serv FECs Recv      : 2
...
-----
At time t = 3 sec (Mode: Absolute)
-----
Addr FECs Sent      : 0                      Addr FECs Recv      : 0
Serv FECs Sent      : 1                      Serv FECs Recv      : 2
...
-----
```

```

At time t = 6 sec (Mode: Absolute)
-----
Addr FECs Sent   : 0                      Addr FECs Recv    : 0
Serv FECs Sent   : 1                      Serv FECs Recv    : 2
...
-----
At time t = 9 sec (Mode: Absolute)
-----
Addr FECs Sent   : 0                      Addr FECs Recv    : 0
Serv FECs Sent   : 1                      Serv FECs Recv    : 2
...
=====
ALU-12>monitor>router>ldp#

```

```

ALU-12>monitor>router>ldp# statistics interval 3 repeat 3 rate
=====
Monitor statistics for LDP instance
=====
At time t = 0 sec (Base Statistics)
-----
Addr FECs Sent   : 0                      Addr FECs Recv    : 0
Serv FECs Sent   : 1                      Serv FECs Recv    : 2
...
-----
At time t = 3 sec (Mode: Rate)
-----
Addr FECs Sent   : 0                      Addr FECs Recv    : 0
Serv FECs Sent   : 0                      Serv FECs Recv    : 0
...
-----
At time t = 6 sec (Mode: Rate)
-----
Addr FECs Sent   : 0                      Addr FECs Recv    : 0
Serv FECs Sent   : 0                      Serv FECs Recv    : 0
...
-----
At time t = 9 sec (Mode: Rate)
-----
Addr FECs Sent   : 0                      Addr FECs Recv    : 0
Serv FECs Sent   : 0                      Serv FECs Recv    : 0
...
=====

```

6.12.2.5 Debug commands

The following output shows debug LDP configurations discussed in this section.

```

ALU-12# debug router ldp peer 10.10.10.104
ALU-12>debug>router>ldp# show debug ldp
debug
  router "Base"
    ldp peer 10.10.10.104
      event
        bindings
        messages
      exit
    packet
      hello
      init
      keepalive

```

```
        label
      exit
    exit
  exit
exit
ALU-12>debug>router>ldp#
```

ldp

Syntax

[no] ldp

Context

debug>router

Description

This command configures LDP debugging.

interface

Syntax

[no] interface *interface-name*

Context

debug>router>ldp

Description

This command configures debugging for a specific LDP interface.

Parameters

interface-name

specifies an existing interface

peer

Syntax

[no] peer *ip-address*

Context

debug>router>ldp

Description

This command configures debugging for a specific LDP peer.

Parameters

ip-address

specifies the LDP peer to debug

event

Syntax

[no] event

Context

debug>router>ldp>interface

debug>router>ldp>peer

Description

This command configures debugging for specific LDP events.

bindings

Syntax

[no] bindings

Context

debug>router>ldp>peer>event

Description

This command displays debugging information about addresses and label bindings learned from LDP peers for LDP bindings.

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

messages

Syntax

[no] messages

Context

debug>router>ldp>if>event

debug>router>ldp>peer>event

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 the command disables debugging output for LDP messages.

packet

Syntax

[no] packet

Context

debug>router>ldp>interface

debug>router>ldp>peer

Description

This command enables debugging for specific LDP packets.

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

hello

Syntax

hello [detail]

no hello

Context

debug>router>ldp>if>packet

debug>router>ldp>peer>packet

Description

This command enables debugging for sent and received LDP Hello packets.

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

Parameters

detail

displays detailed information

init

Syntax

init [detail]

no init**Context**

debug>router>ldp>peer>packet

Description

This command enables debugging for LDP Init packets. The **detail** option displays detailed information about the type length value (TLV) included in mac-flush packets.

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

Parameters**detail**

displays detailed information

keepalive**Syntax**

[no] keepalive

Context

debug>router>ldp>peer>packet

Description

This command enables debugging for LDP keepalive packets.

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

label**Syntax**

label [detail]

no label

Context

debug>router>ldp neighbor>packet

Description

This command enables debugging for LDP label packets.

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

Parameters**detail**

displays detailed information

7 List of acronyms

Table 55: Acronyms

Acronym	Expansion
2G	second-generation wireless telephone technology
3DES	triple DES (data encryption standard)
3G	third-generation mobile telephone technology
6VPE	IPv6 on virtual private edge router
7705 SAR	7705 Service Aggregation Router
7750 SR	7750 Service Router
8 PSK	eight phase shift keying
16 QAM	16-state quadrature amplitude modulation
32 QAM	32-state quadrature amplitude modulation
64 QAM	64-state quadrature amplitude modulation
128 QAM	128-state quadrature amplitude modulation
256 QAM	256-state quadrature amplitude modulation
ABR	area border router available bit rate
AC	alternating current attachment circuit
ACK	acknowledge
ACL	access control list
ACR	adaptive clock recovery
AD	auto-discovery
ADC	asymmetric delay control
ADM	add/drop multiplexer
ADP	automatic discovery protocol
AES	advanced encryption standard

Acronym	Expansion
AFI	authority and format identifier
AIGP	accumulated IGP
AIS	alarm indication signal
ALG	application level gateway
AMP	active multipath
AN	association number
ANSI	American National Standards Institute
Apipe	ATM VLL
APS	automatic protection switching
ARP	address resolution protocol
A/S	active/standby
AS	autonomous system
ASAP	any service, any port
ASBR	autonomous system boundary router
ASM	any-source multicast autonomous system message
ASN	autonomous system number
ATM	asynchronous transfer mode
ATM PVC	ATM permanent virtual circuit
AU	administrative unit
AUG	administrative unit group
B3ZS	bipolar with three-zero substitution
Batt A	battery A
B-bit	beginning bit (first packet of a fragment)
BBE	background block errors
Bc	committed burst size
Be	excess burst size
BECN	backward explicit congestion notification

Acronym	Expansion
Bellcore	Bell Communications Research
BFD	bidirectional forwarding detection
BGP	border gateway protocol
BGP-LS	border gateway protocol link state
BGP-LU	border gateway protocol labeled unicast
BITS	building integrated timing supply
BTCA	best timeTransmitter clock algorithm
BMU	broadcast, multicast, and unknown traffic Traffic that is not unicast. Any nature of multipoint traffic: <ul style="list-style-type: none"> • broadcast (that is, all 1s as the destination IP to represent all destinations within the subnet) • multicast (that is, traffic typically identified by the destination address, uses special destination address); for IP, the destination must be 224.0.0.0 to 239.255.255.255 • unknown (that is, the destination is typically a valid unicast address but the destination port/interface is not yet known; therefore, traffic needs to be forwarded to all destinations; unknown traffic is treated as broadcast)
BNM	bandwidth notification message
BOF	boot options file
BoS	bottom of stack
BPDU	bridge protocol data unit
BRAS	Broadband Remote Access Server
BSC	Base Station Controller
BSM	bootstrap message
BSR	bootstrap router
BSTA	Broadband Service Termination Architecture
BTS	base transceiver station
CA	certificate authority connectivity association
CAK	connectivity association key

Acronym	Expansion
CAS	channel associated signaling
CBN	common bonding networks
CBS	committed buffer space
CC	continuity check control channel
CCM	continuity check message
CCTV	closed-circuit television
CE	circuit emulation customer edge
CEM	circuit emulation
CES	circuit emulation services
CESoPSN	circuit emulation services over packet switched network
CFM	connectivity fault management
cHDLC	Cisco high-level data link control protocol
CIDR	classless inter-domain routing
CIR	committed information rate
CKN	connectivity association key name
CLI	command line interface
CLP	cell loss priority
CMP	certificate management protocol
CMV	Cryptographic Module version
C-multicast	customer multicast
CoS	class of service
CPE	customer premises equipment
Cpipe	circuit emulation (or TDM) VLL
CPM	Control and Processing Module (CPM is used instead of CSM when referring to CSM filtering to align with CLI syntax used with other SR products). CSM management ports are referred to as CPM management ports in the CLI.

Acronym	Expansion
CPROTO	C prototype
CPU	central processing unit
C/R	command/response
CRC	cyclic redundancy check
CRC-32	32-bit cyclic redundancy check
CRL	certificate revocation list
CRON	a time-based scheduling service (from chronos = time)
CRP	candidate RP
CSM	Control and Switching Module
CSNP	complete sequence number PDU
CSPF	constrained shortest path first
C-tag	customer VLAN tag
CV	connection verification customer VLAN (tag)
CW	control word
CWDM	coarse wavelength-division multiplexing
DA/FAN	distribution automation and field area network
DC	direct current
DC-C	DC return - common
DCE	data communications equipment
DC-I	DC return - isolated
DCO	digitally controlled oscillator
DCR	differential clock recovery
DDoS	distributed DoS
DE	discard eligibility
DER	distinguished encoding rules
DES	data encryption standard
DF	do not fragment

Acronym	Expansion
	designated forwarder
DH	Diffie-Hellman
DHB	decimal, hexadecimal, or binary
DHCP	dynamic host configuration protocol
DHCPv6	dynamic host configuration protocol for IPv6
DIS	designated intermediate system
DLCI	data link connection identifier
DLCMI	data link connection management interface
DM	delay measurement
DNS	domain name server
DNU	do not use
DoS	denial of service
dot1p	IEEE 802.1p bits, in Ethernet or VLAN ingress packet headers, used to map traffic to up to eight forwarding classes
dot1q	IEEE 802.1q encapsulation for Ethernet interfaces
DPD	dead peer detection
DPI	deep packet inspection
DPLL	digital phase locked loop
DR	designated router
DS	digital signature
DSA	digital signal algorithm
DSCP	differentiated services code point
DSL	digital subscriber line
DSLAM	digital subscriber line access multiplexer
DTE	data termination equipment
DU	downstream unsolicited
DUID	DHCP unique identifier
DUS	do not use for synchronization

Acronym	Expansion
DV	delay variation
DVMRP	distance vector multicast routing protocol
e911	enhanced 911 service
EAP	Extensible Authentication Protocol
EAPOL	EAP over LAN
E-bit	ending bit (last packet of a fragment)
E-BSR	elected BSR
ECMP	equal cost multipath
EE	end entity
EFM	Ethernet in the first mile
EGP	exterior gateway protocol
EIA/TIA-232	Electronic Industries Alliance/Telecommunications Industry Association Standard 232 (also known as RS-232)
EIR	excess information rate
EJBCA	Enterprise Java Beans Certificate Authority
E-LAN	Ethernet local area network
E-Line	Ethernet virtual private line
EL	entropy label
eLER	egress label edge router
ELI	entropy label indicator
E&M	ear and mouth earth and magneto exchange and multiplexer
eMBMS	evolved MBMS
EOP	end of packet
EPC	evolved packet core
EPD	early packet discard
Epipe	Ethernet VLL

Acronym	Expansion
EPL	Ethernet private line
EPON	Ethernet Passive Optical Network
EPS	equipment protection switching
ERO	explicit route object
ES	Ethernet segment errored seconds
ESD	electrostatic discharge
ESI	Ethernet segment identifier
ESMC	Ethernet synchronization message channel
ESN	extended sequence number
ESP	encapsulating security payload
ESPI	encapsulating security payload identifier
ESV	Entropy Source version
ETE	end-to-end
ETH-BN	Ethernet bandwidth notification
ETH-CFM	Ethernet connectivity fault management (IEEE 802.1ag)
EVC	Ethernet virtual connection
EVDO	evolution - data optimized
EVI	EVPN instance
EVPL	Ethernet virtual private link
EVPN	Ethernet virtual private network
EXP bits	experimental bits (currently known as TC)
FC	forwarding class
FCS	frame check sequence
FD	frequency diversity
FDB	forwarding database
FDL	facilities data link
FEAC	far-end alarm and control

Acronym	Expansion
FEC	forwarding equivalence class
FECN	forward explicit congestion notification
FeGW	far-end gateway
FEP	front-end processor
FF	fixed filter
FFD	fast fault detection
FIB	forwarding information base
FIFO	first in, first out
FIPS	Federal Information Processing Standard
FM	fault management
FNG	fault notification generator
FOM	figure of merit
Fpipe	frame relay VLL
FPV	FIPS Provider version
FQDN	fully qualified domain name
FR	frame relay
FRG bit	fragmentation bit
FRR	fast reroute
FTN	FEC-to-NHLFE
FTP	file transfer protocol
FXO	foreign exchange office
FXS	foreign exchange subscriber
GFP	generic framing procedure
GigE	Gigabit Ethernet
GLONASS	Global Navigation Satellite System (Russia)
GNSS	global navigation satellite system (generic)
GPON	Gigabit Passive Optical Network

Acronym	Expansion
GPRS	general packet radio service
GPS	Global Positioning System
GRE	generic routing encapsulation
GRT	global routing table
GSM	Global System for Mobile Communications (2G)
GTP-U	GPRS tunneling protocol user plane
GW	gateway
HA	high availability
HCM	high capacity multiplexing
HDB3	high density bipolar of order 3
HDLC	high-level data link control protocol
HEC	header error control
HMAC	hash message authentication code
Hpipe	HDLC VLL
H-QoS	hierarchical quality of service
HSB	hot standby
HSDPA	high-speed downlink packet access
HSPA	high-speed packet access
H-VPLS	hierarchical virtual private line service
IANA	Internet Assigned Numbers Authority
IBN	isolated bonding networks
ICB	inter-chassis backup
ICK	integrity connection value key
ICMP	Internet control message protocol
ICMPv6	Internet control message protocol for IPv6
ICP	IMA control protocol cells
ICV	integrity connection value
IDS	intrusion detection system

Acronym	Expansion
IDU	indoor unit
IED	intelligent end device
IEEE	Institute of Electrical and Electronics Engineers
IEEE 1588v2	Institute of Electrical and Electronics Engineers standard 1588-2008
IES	Internet enhanced service
IETF	Internet Engineering Task Force
IGMP	Internet group management protocol
IGP	interior gateway protocol
IID	instance ID
IKE	Internet key exchange
iLER	ingress label edge router
ILM	incoming label map
IMA	inverse multiplexing over ATM
IMET-IR	inclusive multicast Ethernet tag—ingress replication
INVARP	inverse address resolution protocol
IOM	input/output module
IP	Internet protocol
IPCP	Internet protocol control protocol
IPIP	IP in IP
Ipipe	IP interworking VLL
I-PMSI	inclusive PMSI
IPoATM	IP over ATM
IPS	intrusion prevention system
IPSec	Internet protocol security
IR	ingress replication
IRB	integrated routing and bridging
ISA	integrated services adapter

Acronym	Expansion
ISAKMP	Internet security association and key management protocol
IS-IS	Intermediate System-to-Intermediate System
IS-IS-TE	IS-IS-traffic engineering (extensions)
ISO	International Organization for Standardization
IW	interworking
JP	join prune
KEK	key encryption key
KG	key group
LB	loopback
lbf-in	pound force inch
LBM	loopback message
LBO	line buildout
LBR	loopback reply
LCP	link control protocol
LDP	label distribution protocol
LER	label edge router
LFA	loop-free alternate
LFIB	label forwarding information base
LIB	label information base
LLDP	link layer discovery protocol
LLDPDU	link layer discovery protocol data unit
LLF	link loss forwarding
LLID	loopback location ID
LM	loss measurement
LMI	local management interface
LOS	line-of-sight loss of signal
LSA	link-state advertisement

Acronym	Expansion
LSDB	link-state database
LSP	label switched path link-state PDU (for IS-IS)
LSPA	LSP attributes
LSR	label switching router link-state request
LSU	link-state update
LT	linktrace
LTE	long term evolution line termination equipment
LTM	linktrace message
LTN	LSP ID to NHLFE
LTR	link trace reply
MA	maintenance association
MAC	media access control
MACsec	media access control security
MA-ID	maintenance association identifier
MBB	make-before-break
MBGP	multicast BGP multiprotocol BGP multiprotocol extensions for BGP
MBMS	multimedia broadcast multicast service
MBS	maximum buffer space maximum burst size media buffer space
MBSP	mobile backhaul service provider
MCAC	multicast connection admission control
MC-APS	multi-chassis automatic protection switching
MC-MLPPP	multiclass multilink point-to-point protocol

Acronym	Expansion
MCS	multicast server multi-chassis synchronization
MCT	MPT craft terminal
MD	maintenance domain
MD5	message digest version 5 (algorithm)
MDA	media dependent adapter
MDDDB	multidrop data bridge
MDL	maintenance data link
MDT	multicast distribution tree
ME	maintenance entity
MED	multi-exit discriminator
MEF	Metro Ethernet Forum
MEG	maintenance entity group
MEG-ID	maintenance entity group identifier
MEN	Metro Ethernet network
MEP	maintenance association endpoint
MFC	multi-field classification
MHD	multi-homed device
MHF	MIP half function
MHN	multi-homed network
MI	member identifier
MIB	management information base
MI-IS-IS	multi-instance IS-IS
MIR	minimum information rate
MKA	MACsec key agreement
MLD	multicast listener discovery
mLDP	multicast LDP
MLPPP	multilink point-to-point protocol

Acronym	Expansion
mLSP	multicast LSP
MoFRR	multicast-only fast reroute
MP	merge point multilink protocol multipoint
MP-BGP	multiprotocol border gateway protocol
MPLS	multiprotocol label switching
MPLSCP	multiprotocol label switching control protocol
MPP	MPT protection protocol
MPR	see Wavence
MPR-e	Microwave Packet Radio (standalone mode)
MPT-HC V2/9558HC	Microwave Packet Transport, High Capacity version 2
MPT-HLC	Microwave Packet Transport, High-Capacity Long-Haul Cubic (ANSI)
MPT-HQAM	Microwave Packet Transport, High Capacity (MPT-HC-QAM) or Extended Power (MPT-XP-QAM) with 512/1024 QAM
MPT-MC	Microwave Packet Transport, Medium Capacity
MPT-XP	Microwave Packet Transport, High Capacity (very high power version of MPT-HC V2/9558HC)
MRAI	minimum route advertisement interval
MRRU	maximum received reconstructed unit
MRU	maximum receive unit
MSDP	Multicast Source Discovery Protocol
MSDU	MAC Service Data Unit
MSK	master session key
MSO	multi-system operator
MS-PW	multi-segment pseudowire
MSS	maximum segment size Microwave Service Switch

Acronym	Expansion
MTIE	maximum time interval error
MTSO	mobile trunk switching office
MTU	maximum transmission unit multi-tenant unit
M-VPLS	management virtual private line service
MVPN	multicast VPN
MVR	multicast VPLS registration
MW	microwave
MWA	microwave awareness
N·m	newton meter
NAT	network address translation
NAT-T	network address translation traversal
NBMA	non-broadcast multiple access (network)
ND	neighbor discovery
NDcPP	Network Device Protection Profile
NE	network element
NET	network entity title
NFM-P	Network Functions Manager - Packet (formerly 5620 SAM)
NGE	network group encryption
NG-MVPN	next generation MVPN
NH	next hop
NHLFE	next hop label forwarding entry
NHOP	next-hop
NIAP	National Information Assurance Partnership
NLOS	non-line-of-sight
NLPID	network level protocol identifier
NLRI	network layer reachability information
NNHOP	next next-hop

Acronym	Expansion
NNI	network-to-network interface
Node B	similar to BTS but used in 3G networks — term is used in UMTS (3G systems) while BTS is used in GSM (2G systems)
NOC	network operations center
NPAT	network port address translation
NRC-F	Network Resource Controller - Flow
NRC-P	Network Resource Controller - Packet
NRC-T	Network Resource Controller - Transport
NRC-X	Network Resource Controller - Cross Domain
NSAP	network service access point
NSD	Network Services Director
NSP	native service processing Network Services Platform
NSSA	not-so-stubby area
NTP	Network Time Protocol
NTR	network timing reference
OADM	optical add/drop multiplexer
OAM	operations, administration, and maintenance
OAMPDU	OAM protocol data units
OC3	optical carrier level 3
OCSP	online certificate status protocol
ODU	outdoor unit
OIF	outgoing interface
OLT	optical line termination
OMC	optical management console
ONT	optical network terminal
OOB	out-of-band
OPX	off premises extension

Acronym	Expansion
ORF	outbound route filtering
OS	operating system
OSI	Open Systems Interconnection (reference model)
OSINLCP	OSI Network Layer Control Protocol
OSPF	open shortest path first
OSPF-TE	OSPF-traffic engineering (extensions)
OSS	operations support system
OSSP	organization specific slow protocol
OTP	one time password
OWAMP	one-way active measurement protocol
P2MP	point to multipoint
PADI	PPPoE active discovery initiation
PADR	PPPoE active discovery request
PAE	port authentication entities
PSB	path state block
PBO	packet byte offset
PBR	policy-based routing
PBX	private branch exchange
PCAP	packet capture
PCC	path computation client
PCE	path computation element
PCEP	Path Computation Element Communication Protocol
PCM	pulse code modulation
PCP	priority code point
PCR	proprietary clock recovery
PDU	power distribution unit protocol data units
PDV	packet delay variation

Acronym	Expansion
PDVT	packet delay variation tolerance
PE	provider edge router
PEAPv0	protected extensible authentication protocol version 0
PEM	privacy enhanced mail
PFoE	power feed over Ethernet
PFS	perfect forward secrecy
PHB	per-hop behavior
PHP	penultimate hop popping
PHY	physical layer
PIC	prefix independent convergence
PID	protocol ID
PIM SSM	protocol independent multicast—source-specific multicast
PIR	peak information rate
PKCS	public key cryptography standards
PKI	public key infrastructure
PLAR	private line automatic ringdown
PLCP	Physical Layer Convergence Protocol
PLR	point of local repair
PLSP	path LSP
PM	performance monitoring
PMSI	P-multicast service interface
P-multicast	provider multicast
PN	packet number
PoE	power over Ethernet
PoE+	power over Ethernet plus
POH	path overhead
POI	purge originator identification

Acronym	Expansion
PoP	point of presence
POS	packet over SONET
PPP	point-to-point protocol
PPPoE	point-to-point protocol over Ethernet
PPS	pulses per second
PRC	primary reference clock
PRS	primary reference source
PRTC	primary reference time clock
PSE	power sourcing equipment
PSK	pre-shared key
PSN	packet switched network
PSNP	partial sequence number PDU
PTA	PMSI tunnel attribute
PTM	packet transfer mode
PTP	performance transparency protocol Precision Time Protocol
PuTTY	an open-source terminal emulator, serial console, and network file transfer application
PVC	permanent virtual circuit
PVCC	permanent virtual channel connection
PW	pseudowire
PWE	pseudowire emulation
PWE3	pseudowire emulation edge-to-edge
Q.922	ITU-T Q-series Specification 922
QL	quality level
QoS	quality of service
QPSK	quadrature phase shift keying
RADIUS	Remote Authentication Dial In User Service

Acronym	Expansion
RAN	radio access network
RBS	robbed bit signaling
RD	route distinguisher
RDI	remote defect indication
RED	random early discard
RESV	reservation
RIB	routing information base
RIP	routing information protocol
RJ45	registered jack 45
RMON	remote network monitoring
RNC	radio network controller
RP	rendezvous point
RPF RTM	reverse path forwarding RTM
RPS	radio protection switching
RPT	rendezvous-point tree
RR	route reflector
RRO	record route object
RS-232	Recommended Standard 232 (also known as EIA/TIA-232)
RSA	Rivest, Shamir, and Adleman (authors of the RSA encryption algorithm)
RSHG	residential split horizon group
RSTP	rapid spanning tree protocol
RSVP-TE	resource reservation protocol - traffic engineering
RT	receive/transmit
RTC	route target constraint
RTM	routing table manager
RTN	battery return
RTP	real-time protocol

Acronym	Expansion
R&TTE	Radio and Telecommunications Terminal Equipment
RTU	remote terminal unit
RU	rack unit
r-VPLS	routed virtual private LAN service
SA	security association source-active
SAA	service assurance agent
SAFI	subsequent address family identifier
SAK	security association key
SAP	service access point
SAToP	structure-agnostic TDM over packet
SCADA	supervisory control and data acquisition
SC-APS	single-chassis automatic protection switching
SCI	secure channel identifier
SCP	secure copy
SCTP	Stream Control Transmission Protocol
SD	signal degrade space diversity
SDH	synchronous digital hierarchy
SDI	serial data interface
SDN	software defined network
SDP	service destination point
SE	shared explicit
SeGW	secure gateway
SES	severely errored seconds
SETS	synchronous equipment timing source
SF	signal fail
SFP	small form-factor pluggable (transceiver)

Acronym	Expansion
SFTP	SSH file transfer protocol
(S,G)	(source, group)
SGT	self-generated traffic
SHA-1	secure hash algorithm
SHG	split horizon group
SID	Segment ID
SIR	sustained information rate
SL	short length
SLA	service-level agreement
SLARP	serial line address resolution protocol
SLID	subscriber location identifier of a GPON module
SLM	synthetic loss measurement
SNMP	Simple Network Management Protocol
SNPA	subnetwork point of attachment
SNR	signal to noise ratio
SNTP	simple network time protocol
SONET	synchronous optical networking
S-PE	switching provider edge router
SPF	shortest path first
SPI	security parameter index
S-PMSI	selective PMSI
SPT	shortest path tree
SR	Service Router (7750 SR) segment routing
SRGB	segment routing global block
SRLG	shared risk link group
SRP	stateful request parameter
SRRP	subscriber routed redundancy protocol

Acronym	Expansion
SR-ISIS	segment routing IS-IS
SR-MPLS	segment routing with MPLS data plane
SR-OSPF	segment routing OSPF
SR-TE	segment routing traffic engineering
SSH	secure shell
SSM	source-specific multicast synchronization status messaging
SSU	system synchronization unit
S-tag	service VLAN tag
STM	synchronous transport module
STM1	synchronous transport module, level 1
STP	spanning tree protocol
STS	synchronous transport signal
SVC	switched virtual circuit
SVEC	synchronization vector
SYN	synchronize
TACACS+	Terminal Access Controller Access-Control System Plus
TC	traffic class (formerly known as EXP bits)
TCI	tag control information
TCP	transmission control protocol
TCP-AO	TCP Authentication Option
TDA	transmit diversity antenna
TDEV	time deviation
TDM	time division multiplexing
TE	traffic engineering
TEDB	traffic engineering database
TEID	tunnel endpoint identifier
TEP	tunnel endpoint

Acronym	Expansion
TFTP	trivial file transfer protocol
T-LDP	targeted LDP
TLS	transport layer security
TLV	type length value
TM	traffic management
ToD	time of day
ToS	type of service
T-PE	terminating provider edge router
TPID	tag protocol identifier
TPIF	IEEE C37.94 teleprotection interface
TPMR	two-port MAC relay
TPS	transmission protection switching
TSoP	transparent SDH/SONET over packet
TTL	time to live
TTLS	tunneled transport layer security
TTM	tunnel table manager
TU	tributary unit
TUG	tributary unit group
TWAMP	two-way active measurement protocol
U-APS	unidirectional automatic protection switching
UAS	unavailable seconds
UBR	unspecified bit rate
UDP	user datagram protocol
UFD	unidirectional forwarding detection
UMH	upstream multicast hop
UMTS	Universal Mobile Telecommunications System (3G)
UNI	user-to-network interface

Acronym	Expansion
uRPF	unicast reverse path forwarding
V.11	ITU-T V-series Recommendation 11
V.24	ITU-T V-series Recommendation 24
V.35	ITU-T V-series Recommendation 35
VC	virtual circuit
VCB	voice conference bridge
VCC	virtual channel connection
VCCV	virtual circuit connectivity verification
VCI	virtual circuit identifier
VID	VLAN ID
VLAN	virtual LAN
VLL	virtual leased line
VM	virtual machine
VoIP	voice over IP
Vp	peak voltage
VP	virtual path
VPC	virtual path connection
VPI	virtual path identifier
VPLS	virtual private LAN service
VPN	virtual private network
VPRN	virtual private routed network
VPWS	virtual private wire service
VRF	virtual routing and forwarding table
VRRP	virtual router redundancy protocol
V-SAP	virtual service access point
VSE	vendor-specific extension
VSI	virtual switch instance

Acronym	Expansion
VSO	vendor-specific option
VT	virtual trunk virtual tributary
VTG	virtual tributary group
Wavence	formerly 9500 MPR (Microwave Packet Radio)
WCDMA	wideband code division multiple access (transmission protocol used in UMTS networks)
WRED	weighted random early discard
WTR	wait to restore
X.21	ITU-T X-series Recommendation 21
XOR	exclusive-OR
XRO	exclude route object

8 Supported standards and protocols

This chapter lists the 7705 SAR compliance with security and telecom standards, the protocols supported, and proprietary MIBs.

8.1 Security standards

FIPS 140-3—Federal Information Processing Standard publication 140-3, Security Requirements for Cryptographic Modules

8.2 Telecom standards

ANSI/TIA/EIA-232-C—Interface Between Data Terminal Equipment and Data Circuit-Terminating Equipment Employing Serial Binary Data Interchange

IEEE 802.1AB-2016—IEEE Standard for Local and metropolitan area networks - Station and Media Access Control Connectivity Discovery

IEEE 802.1ad—IEEE Standard for Local and Metropolitan Area Networks—Virtual Bridged Local Area Networks

IEEE Std 802.1AE-2006 Media Access Control (MAC) Security

IEEE Std 802.1AEbw-2013—Media Access Control (MAC) Security Amendment 2: Extended Packet Numbering

IEEE 802.1ag—Service Layer OAM

IEEE 802.1p/q—VLAN Tagging

IEEE 802.1x-2010—IEEE Standard for Local and Metropolitan Area Networks—Port-based Network Access Control

IEEE 802.3—10BaseT

IEEE 802.3ab—1000BaseT

IEEE 802.3ah—Ethernet OAM

IEEE 802.3u—100BaseTX

IEEE 802.3x —Flow Control

IEEE 802.3z—1000BaseSX/LX

IEEE 802.3-2008—Revised base standard

IEEE 802.1AX-2008—Link Aggregation Task Force (transferred from IEEE 802.3ad)

IEEE C37.94-2017—N Times 64 Kilobit Per Second Optical Fiber Interfaces Between Teleprotection and Multiplexer Equipment

ITU-T G.704—Synchronous frame structures used at 1544, 6312, 2048, 8448 and 44 736 kbit/s hierarchical levels

ITU-T G.707—Network node interface for the Synchronous Digital Hierarchy (SDH)

ITU-T G.826—End-to-end error performance parameters and objectives for international, constant bit-rate digital paths and connections

ITU-T G.8032 — Ethernet Ring Protection Switching

ITU-T G.984.1—Gigabit-capable passive optical networks (GPON): general characteristics

ITU-T Y.1564—Ethernet service activation test methodology

ITU-T Y.1731—OAM functions and mechanisms for Ethernet-based networks

8.3 Protocol support

8.3.1 ATM

AF-PHY-0086.001—Inverse Multiplexing for ATM (IMA)

af-tm-0121.000—Traffic Management Specification Version 4.1, March 1999

GR-1113-CORE—Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

GR-1248-CORE—Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996

ITU-T Recommendation I.432.1—B-ISDN user-network interface - Physical layer specification: General characteristics

ITU-T Recommendation I.610—B-ISDN Operation and Maintenance Principles and Functions version 11/95

RFC 2514—Definitions of Textual Conventions and OBJECT_IDENTITIES for ATM Management, February 1999

RFC 2515—Definition of Managed Objects for ATM Management, February 1999

RFC 2684—Multiprotocol Encapsulation over ATM Adaptation Layer 5

8.3.2 BFD

RFC 7130—Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces

RFC 7881—Seamless Bidirectional Forwarding Detection (S-BFD) for IPv4, IPv6, and MPLS

draft-ietf-bfd-mib-00.txt—Bidirectional Forwarding Detection Management Information Base

draft-ietf-bfd-base-o5.txt—Bidirectional Forwarding Detection

draft-ietf-bfd-v4v6-1hop-06.txt—BFD IPv4 and IPv6 (Single Hop)

draft-ietf-bfd-multihop-06.txt—BFD for Multi-hop Paths

8.3.3 BGP

RFC 1397—BGP Default Route Advertisement
RFC 1997—BGP Communities Attribute
RFC 2385—Protection of BGP Sessions via the TCP MD5 Signature Option
RFC 2439—BGP Route Flap Dampening
RFC 2545—Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing
RFC 2918—Route Refresh Capability for BGP-4
RFC 3107—Carrying Label Information in BGP-4
RFC 3392—Capabilities Advertisement with BGP-4
RFC 4271—BGP-4 (previously RFC 1771)
RFC 4360—BGP Extended Communities Attribute
RFC 4364—BGP/MPLS IP Virtual Private Networks (VPNs) (previously RFC 2574bis BGP/MPLS VPNs)
RFC 4456—BGP Route Reflection: Alternative to Full-mesh IBGP (previously RFC 1966 and RFC 2796)
RFC 4486—Subcodes for BGP Cease Notification Message
RFC 4684—Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)
RFC 4724—Graceful Restart Mechanism for BGP - GR Helper
RFC 4760—Multi-protocol Extensions for BGP (previously RFC 2858)
RFC 4893—BGP Support for Four-octet AS Number Space
RFC 4798—Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)
RFC 5549—Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop
RFC 5925—The TCP Authentication Option
RFC 5926—Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)
RFC 6513—Multicast in MPLS/BGP IP VPNs
RFC 6514—BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs
RFC 7311—The Accumulated IGP Metric Attribute for BGP
RFC 7606—Revised Error Handling for BGP UPDATE Messages
draft-ietf-idr-add-paths-04.txt—Advertisement of Multiple Paths in BGP
draft-ietf-idr-add-paths-guidelines-00.txt—Best Practices for Advertisement of Multiple Paths in BGP
draft-weis-esp-group-counter-cipher-00.txt—Using Counter Modes with Encapsulating Security Payload (ESP) and Authentication Header (AH) to Protect Group Traffic

8.3.4 DHCP/DHCPv6

RFC 1534—Interoperation between DHCP and BOOTP
RFC 2131—Dynamic Host Configuration Protocol (REV)

RFC 2132—DHCP Options and BOOTP Vendor Extensions
RFC 3046—DHCP Relay Agent Information Option (Option 82)
RFC 3315—Dynamic Host Configuration Protocol for IPv6
RFC 3736—Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6

8.3.5 Differentiated services

RFC 2474—Definition of the DS Field in the IPv4 and IPv6 Headers
RFC 2597—Assured Forwarding PHB Group
RFC 2598—An Expedited Forwarding PHB
RFC 3140—Per-Hop Behavior Identification Codes

8.3.6 Digital data network management

V.35
RS-232 (also known as EIA/TIA-232)
X.21

8.3.7 ECMP

RFC 2992—Analysis of an Equal-Cost Multi-Path Algorithm

8.3.8 Ethernet VPN (EVPN)

RFC 7432—BGP MPLS-Based Ethernet VPN
draft-ietf-bess-evpn-vpws—Virtual Private Wire Service support in Ethernet VPN
draft-ietf-rabadan-bess-evpn-pref-pdf—Preference-based EVPN DF Election

8.3.9 Frame relay

ANSI T1.617 Annex D—Signalling Specification For Frame Relay Bearer Service
ITU-T Q.922 Annex A—Digital Subscriber Signalling System No. 1 (DSS1) data link layer - ISDN data link layer specification for frame mode bearer services
FRF.1.2—PVC User-to-Network Interface (UNI) Implementation Agreement
RFC 2427—Multiprotocol Interconnect over Frame Relay

8.3.10 GRE

RFC 2784—Generic Routing Encapsulation (GRE)

8.3.11 Internet protocol (IP) – version 4

RFC 768—User Datagram Protocol
RFC 791—Internet Protocol
RFC 792—Internet Control Message Protocol
RFC 793—Transmission Control Protocol
RFC 826—Ethernet Address Resolution Protocol
RFC 854—Telnet Protocol Specification
RFC 1350—The TFTP Protocol (Rev. 2)
RFC 1812—Requirements for IPv4 Routers
RFC 3021—Using 31-Bit Prefixes on IPv4 Point-to-Point Links

8.3.12 Internet protocol (IP) – version 6

RFC 2460—Internet Protocol, Version 6 (IPv6) Specification
RFC 2462—IPv6 Stateless Address Autoconfiguration
RFC 2464—Transmission of IPv6 Packets over Ethernet Networks
RFC 3587—IPv6 Global Unicast Address Format
RFC 3595—Textual Conventions for IPv6 Flow Label
RFC 4007—IPv6 Scoped Address Architecture
RFC 4193—Unique Local IPv6 Unicast Addresses
RFC 4291—IPv6 Addressing Architecture
RFC 4443—Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 Specification
RFC 4649—DHCPv6 Relay Agent Remote-ID Option
RFC 4861—Neighbor Discovery for IP version 6 (IPv6)
RFC 5095—Deprecation of Type 0 Routing Headers in IPv6
RFC 5952—A Recommendation for IPv6 Address Text Representation

8.3.13 IPSec

ITU-T X.690 (2002)—ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)
PKCS #12 Personal Information Exchange Syntax Standard
RFC 2315—PKCS #7: Cryptographic Message Syntax
RFC 2409—The Internet Key Exchange (IKE)
RFC 2986—PKCS #10: Certification Request Syntax Specification
RFC 3706—A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers

RFC 3947—Negotiation of NAT-Traversal in the IKE
RFC 3948—UDP Encapsulation of IPsec ESP Packets
RFC 4301—Security Architecture for the Internet Protocol
RFC 4303—IP Encapsulating Security Payload (ESP)
RFC 4210—Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)
RFC 4211—Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)
RFC 4945—The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2, and PKIX
RFC 5280—Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile
RFC 5996—Internet Key Exchange Protocol Version 2 (IKEv2)
RFC 7383—Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation

8.3.14 IS-IS

RFC 1142—OSI IS-IS Intra-domain Routing Protocol (ISO 10589)
RFC 1195—Use of OSI IS-IS for routing in TCP/IP & dual environments
RFC 2763—Dynamic Hostname Exchange for IS-IS
RFC 2966—Domain-wide Prefix Distribution with Two-Level IS-IS
RFC 2973—IS-IS Mesh Groups
RFC 3373—Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies
RFC 3567—Intermediate System to Intermediate System (IS-IS) Cryptographic Authentication
RFC 3719—Recommendations for Interoperable Networks using IS-IS
RFC 3784—Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)
RFC 3787—Recommendations for Interoperable IP Networks
RFC 4205 for Shared Risk Link Group (SRLG) TLV
RFC 4971—Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information
RFC 5120—M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)
RFC 5304—IS-IS Cryptographic Authentication
RFC 5305—IS-IS Extensions for Traffic Engineering
RFC 5307—IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)
RFC 5308—Routing IPv6 with IS-IS
RFC 5309—Point-to-Point Operation over LAN in Link State Routing Protocols
RFC 5310—IS-IS Generic Cryptographic Authentication
RFC 6232—Purge Originator Identification TLV for IS-IS
RFC 8570—IS-IS Traffic Engineering (TE) Metric Extensions

8.3.15 LDP

RFC 5036—LDP Specification

RFC 5283—LDP Extension for Inter-Area Label Switched Paths

RFC 5350—IANA Considerations for the IPv4 and IPv6 Router Alert Options

RFC 5443—LDP IGP Synchronization

RFC 5561—LDP Capabilities

RFC 5926—Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)

RFC 6388—Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths

RFC 6512—Using Multipoint LDP When the Backbone Has No Route to the Root

RFC 6829—Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6

RFC 7552—Updates to LDP for IPv6

draft-ietf-mpls-ldp-ip-pw-capability—Controlling State Advertisements Of Non-negotiated LDP Applications

draft-ietf-mpls-oam-ipv6-rao—IPv6 Router Alert Option for MPLS OAM

draft-pdutta-mpls-ldp-adj-capability-00—LDP Adjacency Capabilities

draft-pdutta-mpls-ldp-v2-00—LDP Version 2

draft-pdutta-mpls-mldp-up-redundancy-00.txt—Upstream LSR Redundancy for Multi-point LDP Tunnels

draft-weis-esp-group-counter-cipher-00—Using Counter Modes with Encapsulating Security Payload (ESP) and Authentication Header (AH) to Protect Group Traffic

8.3.16 LDP and IP FRR

RFC 5286—Basic Specification for IP Fast Reroute: Loop-Free Alternates

RFC 7490—Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)

8.3.17 MPLS

RFC 3031—MPLS Architecture

RFC 3032—MPLS Label Stack Encoding

RFC 3815—Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)

RFC 5440—Path Computation Element (PCE) Communication Protocol (PCEP)

RFC 6790—The Use of Entropy Labels in MPLS Forwarding

RFC 8253—PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)

RFC 8697—Path Computation Element Communication Protocol (PCEP) Extensions for Establishing Relationships between Sets of Label Switched Paths (LSPs)

RFC 8745—Path Computation Element Communication Protocol (PCEP) Extensions for Associating Working and Protection Label Switched Paths (LSPs) with Stateful PCE
RFC 8800—Path Computation Element Communication Protocol (PCEP) Extension for Label Switched Path (LSP) Diversity Constraint Signaling
draft-dhody-pce-pceps-tls13-02—Updates for PCEPS
draft-ietf-pce-stateful-pce—PCEP Extensions for Stateful PCE
draft-ietf-pce-segment-routing—PCEP Extensions for Segment Routing
draft-alvarez-pce-path-profiles—PCE Path Profiles

8.3.18 MPLS – OAM

RFC 6424— Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels
RFC 8029—Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures

8.3.19 Multicast

RFC 3956—Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address
RFC 3973—Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)
RFC 4610—Anycast-RP Using Protocol Independent Multicast (PIM), which is similar to RFC 3446—Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)
RFC 6514—BGP Encodings and Procedures for Multicast in MPLS/IP VPNs
RFC 6826—Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths
cisco-ipmulticast/pim-autorp-spec—Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast, which is similar to RFC 5059—Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)
draft-ietf-l2vpn-vpls-pim-snooping-07—Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)
draft-ietf-mboned-msdp-deploy-nn.txt—Multicast Source Discovery Protocol (MSDP) Deployment Scenarios

8.3.20 Network management

IANA-IFType-MIB
ITU-T X.721—Information technology- OSI-Structure of Management Information
ITU-T X.734—Information technology- OSI-Systems Management: Event Report Management Function
M.3100/3120—Equipment and Connection Models
RFC 1157—SNMPv1
RFC 1850—OSPF-MIB

RFC 1907—SNMPv2-MIB
RFC 2011—IP-MIB
RFC 2012—TCP-MIB
RFC 2013—UDP-MIB
RFC 2030—Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI
RFC 2096—IP-FORWARD-MIB
RFC 2138—RADIUS
RFC 2206—RSVP-MIB
RFC 2571—SNMP-FRAMEWORKMIB
RFC 2572—SNMP-MPD-MIB
RFC 2573—SNMP-TARGET-&-NOTIFICATION-MIB
RFC 2574—SNMP-USER-BASED-SMMIB
RFC 2575—SNMP-VIEW-BASED ACM-MIB
RFC 2576—SNMP-COMMUNITY-MIB
RFC 2588—SONET-MIB
RFC 2665—EtherLike-MIB
RFC 2819—RMON-MIB
RFC 2863—IF-MIB
RFC 2864—INVERTED-STACK-MIB
RFC 3014—NOTIFICATION-LOG MIB
RFC 3164—The BSD Syslog Protocol
RFC 3273—HCRMON-MIB
RFC 3411—An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks
RFC 3412—Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
RFC 3413—Simple Network Management Protocol (SNMP) Applications
RFC 3414—User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)
RFC 3418—SNMP MIB
RFC 3954—Cisco Systems NetFlow Services Export Version 9
RFC 5101—Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information
RFC 5102—Information Model for IP Flow Information Export
draft-ietf-disman-alarm-mib-04.txt
draft-ietf-mpls-ldp-mib-07.txt
draft-ietf-ospf-mib-update-04.txt
draft-ietf-mpls-lsr-mib-06.txt

draft-ietf-mpls-te-mib-04.txt

TMF 509/613—Network Connectivity Model

8.3.21 OSPF

RFC 1765—OSPF Database Overflow

RFC 2328—OSPF Version 2

RFC 2370—Opaque LSA Support

RFC 2740—OSPF for IPv6

RFC 3101—OSPF NSSA Option

RFC 3137—OSPF Stub Router Advertisement

RFC 3509—Alternative Implementations of OSPF Area Border Routers

RFC 3623—Graceful OSPF Restart (support for Helper mode)

RFC 3630—Traffic Engineering (TE) Extensions to OSPF

RFC 4203 for Shared Risk Link Group (SRLG) sub-TLV

RFC 4577—OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs) (support for basic OSPF at PE-CE links)

RFC 4915—Multi-Topology (MT) Routing in OSPF

RFC 4970—Extensions to OSPF for Advertising Optional Router Capabilities

RFC 5185—OSPF Multi-Area Adjacency

8.3.22 OSPFv3

RFC 4552—Authentication/Confidentiality for OSPFv3

8.3.23 PPP

RFC 1332—PPP Internet Protocol Control Protocol (IPCP)

RFC 1570—PPP LCP Extensions

RFC 1619—PPP over SONET/SDH

RFC 1661—The Point-to-Point Protocol (PPP)

RFC 1662—PPP in HDLC-like Framing

RFC 1989—PPP Link Quality Monitoring

RFC 1990—The PPP Multilink Protocol (MP)

RFC 2686—The Multi-Class Extension to Multi-Link PPP

8.3.24 Pseudowires

Metro Ethernet Forum—Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks

RFC 3550—RTP: A Transport Protocol for Real-Time Applications

RFC 3985—Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture

RFC 4385—Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN

RFC 4446—IANA Allocation for PWE3

RFC 4447—Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)

RFC 4448—Encapsulation Methods for Transport of Ethernet over MPLS Networks

RFC 4553—Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)

RFC 4717—Encapsulation Methods for Transport of Asynchronous Transfer Mode (ATM) over MPLS Networks

RFC 4618—Encapsulation Methods for Transport of PPP/High-Level Data Link Control (HDLC) over MPLS Networks

RFC 4619—Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks

RFC 4816—Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service

RFC 5085—Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires

RFC 5086—Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)

draft-ietf-pwe3-redundancy-02.txt—Pseudowire (PW) Redundancy

8.3.25 RIP

RFC 1058—Routing Information Protocol

RFC 2453—RIP Version 2

8.3.26 RADIUS

RFC 2865—Remote Authentication Dial In User Service

RFC 2866—RADIUS Accounting

RFC 6613—RADIUS over TCP

RFC 6614—Transport Layer Security (TLS) Encryption for RADIUS

8.3.27 RSVP-TE and FRR

RFC 2430—A Provider Architecture for DiffServ & TE

RFC 2702—Requirements for Traffic Engineering over MPLS
RFC 2747—RSVP Cryptographic Authentication
RFC 2961—RSVP Refresh Overhead Reduction Extensions
RFC 3097—RSVP Cryptographic Authentication - Updated Message Type Value
RFC 3209—Extensions to RSVP for LSP Tunnels
RFC 3210—Applicability Statement for Extensions to RSVP for LSP Tunnels
RFC 3477—Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)
RFC 4090—Fast Reroute Extensions to RSVP-TE for LSP Tunnels

8.3.28 Segment routing (SR)

draft-francois-rtgwg-segment-routing-ti-lfa-04—Topology Independent Fast Reroute using Segment Routing
draft-gredler-idr-bgp-ls-segment-routing-ext-03—BGP Link-State extensions for Segment Routing
draft-ietf-isis-segment-routing-extensions-04—IS-IS Extensions for Segment Routing
draft-ietf-mpls-spring-lsp-ping-02—Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane
draft-ietf-ospf-segment-routing-extensions-04—OSPF Extensions for Segment Routing
draft-ietf-spring-segment-routing-15—Segment Routing Architecture

8.3.29 SONET/SDH

GR-253-CORE—SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000
ITU-T Recommendation G.841—Telecommunication Standardization Section of ITU, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum1 issued in July 2002

8.3.30 SSH

RFC 4253—The Secure Shell (SSH) Transport Layer Protocol
draft-ietf-secsh-architecture.txt—SSH Protocol Architecture
draft-ietf-secsh-userauth.txt—SSH Authentication Protocol
draft-ietf-secsh-connection.txt—SSH Connection Protocol
draft-ietf-secsh-newmodes.txt—SSH Transport Layer Encryption Modes
draft-ietf-secsh-filexfer-13.txt—SSH File Transfer Protocol

8.3.31 Synchronization

G.781—Synchronization layer functions, 2001/09/17

G.803—Architecture of transport networks based on the synchronous digital hierarchy (SDH)

G.813—Timing characteristics of SDH equipment slave clocks (SEC)

G.823—The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy, 2003/03/16

G.824—The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy, 2003/03/16

G.8261—Timing and synchronization aspects in packet networks

G.8262—Timing characteristics of synchronous Ethernet equipment slave clock

GR 1244 CORE—Clocks for the Synchronized Network: Common Generic Criteria

IEC/IEEE 61850-9-3—Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation

IEEE C37.238-2017 - IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications

IEEE Std 1588-2008—IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems

IEEE Std 1588-2008—IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, Annex E – Transport of PTP over User Datagram Protocol over Internet Protocol Version 6

IEEE Std 1588-2019—IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, Annex J

ITU-T G.8264—Telecommunication Standardization Section of ITU, Distribution of timing information through packet networks, issued 10/2008

ITU-T G.8265.1—Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for frequency synchronization, issued 10/2010

ITU-T G.8275.1—Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014

ITU-T G.8275.2—Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for time/phase synchronization with partial timing support from the network, issued 06/2016

RFC 5905—Network Time Protocol Version 4: Protocol and Algorithms Specification

RFC 8573—Message Authentication Code for the Network Time Protocol

8.3.32 TACACS+

IETF draft-grant-tacacs-02.txt—The TACACS+ Protocol

8.3.33 TLS

RFC 5246—The Transport Layer Security (TLS) Protocol Version 1.2

RFC 5425—Transport Layer Security (TLS) Transport Mapping for Syslog

RFC 5922—Domain Certificates in the Session Initiation Protocol (SIP)

RFC 6460—Suite B Profile for Transport Layer Security (TLS)

RFC 8446—The Transport Layer Security (TLS) Protocol Version 1.3

8.3.34 TWAMP

RFC 5357—A Two-Way Active Measurement Protocol (TWAMP)

RFC 8762—Simple Two-Way Active Measurement Protocol

RFC 8972—Simple Two-Way Active Measurement Protocol Optional Extensions

8.3.35 VPLS

RFC 4762—Virtual Private LAN Services Using LDP

8.3.36 VRRP

RFC 2787—Definitions of Managed Objects for the Virtual Router Redundancy Protocol

RFC 3768 Virtual Router Redundancy Protocol

RFC 5798 Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6

8.4 Proprietary MIBs

TIMETRA-ATM-MIB.mib

TIMETRA-CAPABILITY-7705-V1.mib

TIMETRA-CHASSIS-MIB.mib

TIMETRA-CLEAR-MIB.mib

TIMETRA-FILTER-MIB.mib

TIMETRA-GLOBAL-MIB.mib

TIMETRA-LAG-MIB.mib

TIMETRA-LDP-MIB.mib

TIMETRA-LOG-MIB.mib

TIMETRA-MPLS-MIB.mib

TIMETRA-OAM-TEST-MIB.mib

TIMETRA-PORT-MIB.mib

TIMETRA-PPP-MIB.mib

TIMETRA-QOS-MIB.mib

TIMETRA-ROUTE-POLICY-MIB.mib

TIMETRA-RSVP-MIB.mib

TIMETRA-SAP-MIB.mib

TIMETRA-SDP-MIB.mib

TIMETRA-SECURITY-MIB.mib


TIMETRA-SERV-MIB.mib

TIMETRA-SYSTEM-MIB.mib

TIMETRA-TC-MIB.mib

TIMETRA-VRRP-MIB.mib

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)