



**7450 ETHERNET SERVICE SWITCH
7750 SERVICE ROUTER
7950 EXTENSIBLE ROUTING SYSTEM**

**MPLS GUIDE
RELEASE 14.0.R4**

3HE 10790 AAAB TQZZA 01

Issue: 01

July 2016

Nokia is a registered trademark of Nokia Corporation. Other products and company names mentioned herein may be trademarks or tradenames of their respective owners.

The information presented is subject to change without notice. No responsibility is assumed for inaccuracies contained herein.

© 2016 Nokia. All rights reserved.

Contains proprietary/trade secret information which is the property of Nokia and must not be made available to, or copied or used by anyone outside Nokia without its written authorization. Not to be used or disclosed except in accordance with applicable agreements.

Table of Contents

1	Getting Started	15
1.1	About This Guide	15
1.2	Nokia Router Configuration Process	16
2	MPLS and RSVP	17
2.1	In This Chapter	17
2.2	MPLS	18
2.2.1	MPLS Label Stack	19
2.2.1.1	Label Values	20
2.2.2	MPLS Entropy Label and Hash Label	22
2.2.2.1	Inserting and Processing the Entropy Label	24
2.2.2.2	Egress LER	25
2.2.2.3	Mapping Entropy Label and Entropy Label Capability at LSP Stitching Points	26
2.2.2.4	Entropy Label on OAM Packets	26
2.2.3	Label Switching Routers	26
2.2.3.1	LSP Types	27
2.2.4	Bidirectional Forwarding Detection for MPLS LSPs	28
2.2.4.1	Bootstrapping the BFD Session	29
2.2.4.2	LSP BFD Configuration	29
2.2.4.3	Enabling and Implementing Limits for LSP BFD on a Node	30
2.2.4.4	BFD Configuration on RSVP-TE LSPs	31
2.2.4.5	MPLS Facility Bypass Method of MPLS Fast Re-Route (FRR)	32
2.2.4.6	Manual Bypass LSP	33
2.2.4.7	Uniform FRR Failover Time	36
2.2.4.8	MPLS/RSVP on Broadcast Interface	37
2.2.4.9	Automatic Bandwidth Allocation for RSVP LSPs	37
2.3	RSVP	47
2.3.1	Using RSVP for MPLS	49
2.3.1.1	RSVP Traffic Engineering Extensions for MPLS	49
2.3.1.2	Hello Protocol	50
2.3.1.3	MD5 Authentication of RSVP Interface	50
2.3.1.4	Configuring Authentication using Keychains	51
2.3.2	Reservation Styles	52
2.3.2.1	RSVP Message Pacing	53
2.3.3	RSVP Overhead Refresh Reduction	53
2.3.4	RSVP Graceful Restart Helper	54
2.3.5	Enhancements to RSVP Control Plane Congestion Control	55
2.3.6	RSVP LSP Statistics	56
2.3.7	P2MP RSVP-TE LSP Statistics	56
2.3.7.1	Configuring RSVP P2MP LSP Egress Statistics	57
2.3.7.2	Configuring RSVP P2MP LSP Ingress Statistics	58
2.3.7.3	Configuring Implicit Null	60
2.3.8	Using Unnumbered Point-to-Point Interface in RSVP	61

2.3.8.1	Operation of RSVP FRR Facility Backup over Unnumbered Interface.....	63
2.4	MPLS Transport Profile	64
2.4.1	MPLS-TP Model	65
2.4.2	MPLS-TP Provider Edge and Gateway	66
2.4.2.1	VLL Services	66
2.4.2.2	Spoke-SDP Termination.....	67
2.4.3	MPLS-TP LSR.....	69
2.4.4	Detailed Descriptions of MPLS-TP	69
2.4.4.1	MPLS-TP LSPs	69
2.4.4.2	MPLS-TP on Pseudowires	70
2.4.5	MPLS-TP Maintenance Identifiers.....	70
2.4.5.1	Generic Associated Channel.....	75
2.4.5.2	MPLS-TP Operations, Administration and Maintenance (OAM)	76
2.4.5.3	PW Control Channel Status Notifications (Static Pseudowire Status Signaling)	80
2.4.5.4	PW Control Channel Status Request Mechanism.....	82
2.4.5.5	Pseudowire Redundancy and Active / Standby Dual Homing	82
2.4.5.6	Lock Instruct and Loopback for MPLS-TP Pseudowires	83
2.4.5.7	MPLS-TP LSP Protection	84
2.4.6	Switching Static MPLS-TP to Dynamic T-LDP Signaled PWs.....	87
2.4.7	Alarm Indication Signal (AIS).....	88
2.4.8	Configuring MPLS-TP	90
2.4.8.1	Configuration Overview	90
2.4.8.2	Node-Wide MPLS-TP Parameter Configuration.....	91
2.4.8.3	Node-Wide MPLS-TP Identifier Configuration.....	92
2.4.8.4	Static LSP and Pseudowire (VC) Label and Tunnel Ranges	92
2.4.8.5	Interface Configuration for MPLS-TP	94
2.4.8.6	LER Configuration for MPLS-TP	96
2.4.8.7	Intermediate LSR Configuration for MPLS-TP LSPs.....	103
2.4.9	MPLS-TP Show Commands.....	104
2.4.9.1	Static MPLS Labels	104
2.4.9.2	MPLS-TP Tunnel Configuration.....	105
2.4.9.3	MPLS-TP Path configuration	106
2.4.9.4	MPLS-TP Protection.....	109
2.4.9.5	MPLS TP Node Configuration	110
2.4.9.6	MPLS-TP Interfaces	112
2.4.10	MPLS-TP Debug Commands	112
2.5	Traffic Engineering	114
2.5.1	TE Metric (IS-IS and OSPF).....	115
2.5.2	Admin Group Support on Facility Bypass Backup LSP	116
2.5.2.1	Procedures at Head-End Node	116
2.5.2.2	Procedures at PLR Node.....	117
2.5.3	Diff-Serv Traffic Engineering	119
2.5.3.1	Mapping of Traffic to a Diff-Serv LSP	119
2.5.3.2	Admission Control of Classes.....	119
2.5.3.3	RSVP Control Plane Extensions	124
2.5.3.4	IGP Extensions.....	124
2.5.3.5	Diff-Serv TE Configuration and Operation.....	125

2.5.4	Diff-Serv TE LSP Class Type Change under Failure	128
2.5.4.1	LSP Primary Path Retry Procedures	129
2.5.4.2	Bandwidth Sharing Across Class Types.....	131
2.5.4.3	Downgrading the CT of Bandwidth Sharing LSP Paths	132
2.5.4.4	Upgrading the CT of Bandwidth Sharing LSP Paths	133
2.6	Advanced MPLS/RSVP Features.....	134
2.6.1	Extending RSVP LSP to use Loopback Interfaces Other than Router-id.....	135
2.6.2	LSP Path Change.....	135
2.6.3	Manual LSP Path Switch.....	136
2.6.4	Make-Before-Break (MBB) Procedures for LSP/Path Parameter Configuration Change.....	137
2.6.5	Automatic Creation of RSVP-TE LSP Mesh.....	138
2.6.6	RSVP-TE LSP Shortcut for IGP Resolution	140
2.6.6.1	Using LSP Relative Metric with IGP Shortcut.....	142
2.6.6.2	ECMP Considerations	143
2.6.6.3	Handling of Control Packets	144
2.6.6.4	Forwarding Adjacency	144
2.6.6.5	LDP Forwarding over IGP Shortcut	145
2.6.6.6	Handling of Multicast Packets	145
2.6.7	RSVP-TE LSP Signaling using LSP Template	147
2.6.8	Shared Risk Link Groups.....	148
2.6.8.1	Enabling Disjoint Backup Paths.....	149
2.6.8.2	SRLG Penalty Weights for Detour and Bypass LSPs	150
2.6.8.3	Static Configurations of SRLG Memberships	152
2.6.9	TE Graceful Shutdown	154
2.6.10	Soft Preemption of Diff-Serv RSVP LSP	154
2.6.11	Least-Fill Bandwidth Rule in CSPF ECMP Selection	154
2.6.12	Inter-Area TE LSP (ERO Expansion Method)	155
2.6.12.1	Area Border Node FRR Protection for Inter-Area LSP	155
2.6.12.2	Inter-Area LSP support of OSPF Virtual Links	159
2.6.12.3	Area Border Node FRR Protection for Inter-Area LSP	159
2.6.13	Automatic Creation of a RSVP Mesh LSP	161
2.6.13.1	Feature Configuration.....	161
2.6.13.2	Feature Behavior	162
2.6.13.3	Multi-Area and Multi-Instance Support	163
2.6.13.4	Mesh LSP Name Encoding and Statistics	164
2.6.14	Timer-based Reversion for RSVP-TE LSPs.....	165
2.6.15	Automatic Creation of an RSVP One-Hop LSP.....	166
2.6.15.1	Feature Configuration.....	166
2.6.15.2	Feature Behavior	166
2.6.16	MPLS Entropy Label	167
2.7	Point-to-Multipoint (P2MP) RSVP LSP.....	167
2.7.1	Application in Video Broadcast.....	168
2.7.2	P2MP LSP Data Plane	168
2.7.2.1	Procedures at Ingress LER Node.....	169
2.7.2.2	Procedures at LSR Node.....	169
2.7.2.3	Procedures at Branch LSR Node	170
2.7.2.4	Procedures at Egress LER Node	170

2.7.2.5	Procedures at BUD LSR Node	170
2.7.3	Ingress Path Management for P2MP LSP Packets	171
2.7.3.1	Ingress P2MP Path Management on XCM/IOM-3/IMMs	172
2.7.3.2	Ingress P2MP Path Management on IOM-2.....	173
2.7.4	RSVP Control Plane in a P2MP LSP.....	174
2.7.5	Forwarding Multicast Packets over RSVP P2MP LSP in the Base Router.....	178
2.7.5.1	Procedures at Ingress LER Node.....	178
2.7.5.2	Procedures at Egress LER Node	179
2.8	Segment Routing With Traffic Engineering (SR-TE)	180
2.8.1	SR-TE Support	181
2.8.2	SR-TE LSP Instantiation.....	183
2.8.2.1	PCC-Initiated and PCC-Controlled LSP	185
2.8.2.2	PCC-Initiated and PCE-Computed/Controlled LSP	189
2.8.3	SR-TE LSP Path Computation	191
2.8.3.1	Service and Shortcut Application SR-TE Label Stack Check	193
2.8.4	SR-TE LSP Protection.....	195
2.8.5	Static Route Resolution using SR-TE LSP.....	197
2.8.6	BGP Shortcuts Using SR-TE LSP	198
2.8.7	BGP Label Route Resolution Using SR-TE LSP.....	198
2.8.8	Service Packet Forwarding using SR-TE LSP	199
2.8.9	Data Path Support	200
2.8.9.1	SR-TE LSP Metric and MTU Settings	203
2.8.9.2	Hash Label and Entropy Label Support.....	204
2.9	MPLS Service Usage	205
2.9.1	Service Distribution Paths	205
2.10	MPLS/RSVP Configuration Process Overview.....	205
2.11	Configuration Notes.....	206
2.12	Configuring MPLS and RSVP with CLI.....	207
2.13	MPLS Configuration Overview	208
2.13.1	LSPs.....	208
2.13.2	Paths	208
2.13.3	Router Interface.....	208
2.13.4	Choosing the Signaling Protocol	209
2.14	Basic MPLS Configuration.....	209
2.15	Common Configuration Tasks	210
2.15.1	Configuring MPLS Components	211
2.15.2	Configuring Global MPLS Parameters	211
2.15.3	Configuring an MPLS Interface	212
2.15.4	Configuring MPLS Paths	213
2.15.5	Configuring an MPLS LSP.....	213
2.15.5.1	Configuring a Static LSP	214
2.15.6	Configuring Manual Bypass Tunnels.....	215
2.16	Configuring and Operating SR-TE.....	216
2.16.1	SR-TE Configuration Prerequisites	216
2.16.2	SR-TE LSP Configuration Overview.....	218
2.16.3	Configuring Path Computation and Control for SR-TE LSP	218
2.16.3.1	Configuring Path Profile and Group for PCC-Initiated and PCE-Computed/Controlled LSP	219

2.16.4	Configuring SR-TE LSP Label Stack Size	219
2.16.5	Configuring Adjacency SID Parameters	219
2.16.6	Configuring PCC-controlled, PCE-computed, and PCE-controlled SR-TE LSPs	220
2.17	Configuring RSVP Parameters	222
2.17.1	Configuring RSVP Message Pacing Parameters	223
2.17.2	Configuring Graceful Shutdown.....	223
2.18	MPLS Configuration Management Tasks.....	224
2.18.1	Deleting MPLS.....	224
2.18.2	Modifying MPLS Parameters.....	224
2.18.3	Modifying an MPLS LSP.....	224
2.18.4	Modifying MPLS Path Parameters	225
2.18.5	Modifying MPLS Static LSP Parameters	226
2.18.6	Deleting an MPLS Interface.....	226
2.19	RSVP Configuration Management Tasks.....	227
2.19.1	Modifying RSVP Parameters.....	227
2.19.2	Modifying RSVP Message Pacing Parameters	227
2.19.3	Deleting an Interface from RSVP	228
2.20	MPLS/RSVP Configuration Command Reference	229
2.20.1	Command Hierarchies.....	229
2.20.1.1	MPLS Commands	229
2.20.1.2	MPLS-TP Commands.....	231
2.20.1.3	LSP Commands	232
2.20.1.4	Isp-bfd Commands	237
2.20.1.5	MPLS Path Commands	237
2.20.1.6	RSVP Commands	237
2.20.2	Command Descriptions	239
2.20.2.1	MPLS Commands	239
2.20.2.2	RSVP Commands	336
2.21	MPLS/RSVP Show, Tools, Router, Clear, and Debug Command Reference	361
2.21.1	Command Hierarchies.....	361
2.21.1.1	Show Commands	361
2.21.1.2	Tools Commands	363
2.21.1.3	Router Commands	364
2.21.1.4	Clear Commands.....	364
2.21.1.5	Debug Commands.....	364
2.21.2	Command Descriptions	366
2.21.2.1	Show Commands	366
2.21.2.2	Show Router BFD session Commands	436
2.21.2.3	Show RSVP Commands	437
2.21.2.4	Tools Commands	457
2.21.2.5	Clear Commands.....	471
2.21.2.6	Debug Commands.....	475
3	GMPLS	483
3.1	In This Chapter	483
3.2	GMPLS	483
3.2.1	Example Applications	484

3.2.1.1	Use Case 1: Dynamic Connection Setup with Constraints.....	484
3.2.1.2	Use Case 2: Multi-Layer Resiliency.....	485
3.3	GMPLS UNI Architecture.....	486
3.3.1	Addressing and End-to-End gLSP Architecture	488
3.4	1830 PSS Identifiers.....	489
3.5	Recovery Reference Models	490
3.5.1	End to End Recovery (IP-layer).....	491
3.5.2	End to End ECMP	491
3.5.3	End to End Load Sharing Using a Load Sharing GMPLS Tunnel Group.....	492
3.5.4	End to End Recovery (GMPLS Layer).....	493
3.5.4.1	Unprotected gLSP	493
3.5.4.2	Full LSP Rerouting	494
3.5.4.3	1: N Protection.....	495
3.5.4.4	Optical Segment Recovery.....	497
3.6	Configuring GMPLS with CLI.....	499
3.7	GMPLS Configuration Overview.....	499
3.8	LMP and IPCC Configuration	500
3.8.1	Configuration of IP Communication Channels for LMP and RSVP	500
3.8.2	Configuring LMP.....	501
3.8.3	Configuring Traffic Engineering Links and Data Bearers	503
3.9	Configuring MPLS Paths for GMPLS	505
3.10	Configuring RSVP in GMPLS.....	506
3.11	Configuring a GMPLS LSP on the UNI.....	508
3.11.1	gLSP Constraints.....	510
3.12	Bandwidth.....	510
3.13	Shared Risk Link Groups.....	511
3.14	Optical Network Segment Recovery.....	513
3.15	Configuration of End-to-End GMPLS Recovery	514
3.16	GMPLS Tunnel Groups	516
3.17	Configuring IP and MPLS in an Overlay Network to Use a GMPLS LSP.....	519
3.18	Configuration Notes.....	519
3.19	GMPLS Configuration Command Reference	521
3.19.1	Command Hierarchies.....	521
3.19.1.1	LMP Commands.....	521
3.19.1.2	GMPLS Commands.....	522
3.19.1.3	GMPLS Tunnel Group Commands.....	523
3.19.2	Command Descriptions	523
3.19.2.1	LMP Commands.....	523
3.19.2.2	GMPLS Commands.....	531
3.20	Show, Clear, Debug, and Tools Command Reference	551
3.20.1	Command Hierarchies.....	551
3.20.1.1	Show Commands	551
3.20.1.2	Clear Commands.....	551
3.20.1.3	Debug Commands.....	552
3.20.1.4	Tools Commands	552
3.20.2	Command Descriptions	553
3.20.2.1	Show Commands	553

3.2.0.2.2	Clear Commands.....	560
3.2.0.2.3	Debug Commands.....	560
3.2.0.2.4	Tools Commands	566
4	PCEP	581
4.1	In This Chapter	581
4.2	Introduction to the Path Computation Element Protocol (PCEP)	581
4.2.1	PCC and PCE Configuration	584
4.2.2	Base Implementation of Path Computation Elements (PCE)	585
4.2.3	PCEP Session Establishment and Maintenance.....	587
4.2.4	PCEP Parameters	589
4.2.4.1	Stateful PCE	589
4.2.4.2	PCEP Extensions in Support of SR-TE LSPs	592
4.2.4.3	LSP Initiation	593
4.3	PCEP Configuration Command Reference	599
4.3.1	Command Hierarchies.....	599
4.3.1.1	PCEP Commands	599
4.3.2	Command Descriptions	600
4.3.2.1	PCEP Commands	600
5	Label Distribution Protocol	605
5.1	In This Chapter	605
5.2	Label Distribution Protocol.....	606
5.2.1	LDP and MPLS.....	606
5.2.2	LDP Architecture	607
5.2.3	Subsystem Interrelationships	608
5.2.3.1	Memory Manager and LDP	609
5.2.3.2	Label Manager.....	609
5.2.3.3	LDP Configuration	609
5.2.3.4	Logger	609
5.2.3.5	Service Manager	610
5.2.4	Execution Flow	610
5.2.4.1	Initialization.....	610
5.2.4.2	Session Lifetime	610
5.2.5	Label Exchange	612
5.2.5.1	Other Reasons for Label Actions.....	612
5.2.5.2	Cleanup	612
5.2.5.3	Configuring Implicit Null Label.....	612
5.2.6	Global LDP Filters	613
5.2.6.1	Per LDP Peer FEC Import and Export Policies	613
5.2.7	Configuring Multiple LDP LSR ID	614
5.2.8	T-LDP hello reduction.....	615
5.2.9	Tracking a T-LDP Peer with BFD	616
5.2.10	Link LDP Hello Adjacency Tracking with BFD.....	616
5.2.11	LDP LSP Statistics	617
5.2.12	MPLS Entropy Label.....	617
5.3	TTL Security for BGP and LDP	617
5.4	ECMP Support for LDP	618
5.4.1	Label Operations	619

5.5	Unnumbered Interface Support in LDP	619
5.5.1	Feature Configuration.....	620
5.5.2	Operation of LDP over an Unnumbered IP Interface.....	620
5.5.2.1	Link LDP	621
5.5.2.2	Targeted LDP	621
5.5.2.3	FEC Resolution	622
5.6	LDP over RSVP Tunnels	622
5.6.1	Signaling and Operation	624
5.6.1.1	LDP Label Distribution and FEC Resolution.....	625
5.6.1.2	Default FEC Resolution Procedure	625
5.6.1.3	FEC Resolution Procedure When prefer-tunnel-in-tunnel is Enabled	626
5.6.2	Rerouting Around Failures.....	626
5.6.2.1	LDP-over-RSVP Tunnel Protection	626
5.6.2.2	ABR Protection	627
5.7	LDP over RSVP Without Area Boundary.....	627
5.7.1	LDP over RSVP and ECMP	628
5.8	Class-Based Forwarding of LDP Prefix Packets over IGP Shortcuts	629
5.8.1	Configuration and Operation	629
5.9	LDP ECMP Uniform Failover.....	631
5.10	LDP Fast-Reroute for IS-IS and OSPF Prefixes.....	633
5.10.1	LDP FRR Configuration.....	633
5.10.1.1	Reducing the Scope of the LFA Calculation by SPF.....	634
5.10.2	LDP FRR Procedures.....	635
5.10.2.1	ECMP Considerations	636
5.10.2.2	LDP FRR and LDP Shortcut.....	636
5.10.2.3	LDP FRR and LDP-over-RSVP	637
5.10.2.4	LDP FRR and RSVP Shortcut (IGP Shortcut)	637
5.10.3	IS-IS and OSPF Support for Loop-Free Alternate Calculation	638
5.10.3.1	Loop-Free Alternate Calculation in the Presence of IGP shortcuts	641
5.10.3.2	Loop-Free Alternate Calculation for Inter-Area/inter-Level Prefixes.....	642
5.10.3.3	Loop-Free Alternate Shortest Path First (LFA SPF) Policies	642
5.11	mLDP Fast Upstream Switchover	643
5.12	LDP FEC to BGP Label Route Stitching.....	643
5.12.1	Configuration	645
5.12.2	Detailed LDP FEC Resolution	646
5.12.3	Detailed BGP Labeled Route Resolution	646
5.12.4	Data Plane Forwarding.....	647
5.13	LDP-SR Stitching for IPv4 prefixes (IS-IS)	647
5.13.1	LDP-SR Stitching Configuration	648
5.13.2	Stitching in the LDP-to-SR Direction	649
5.13.3	Stitching in the SR-to-LDP Direction	651
5.14	LDP FRR Remote LFA Backup using SR Tunnel for IPv4 Prefixes (IS-IS)	653
5.15	Automatic Creation of a Targeted Hello Adjacency and LDP Session.....	654
5.15.1	Feature Configuration.....	654
5.15.2	Feature Behavior	656

5.16	Multicast P2MP LDP for GRT	658
5.17	LDP P2MP Support	659
5.17.1	LDP P2MP Configuration	659
5.17.2	LDP P2MP Protocol.....	660
5.17.3	Make Before Break (MBB).....	660
5.17.4	ECMP Support.....	660
5.17.5	Inter-AS Non-segmented mLDP	661
5.17.5.1	In-band Signaling with Non-segmented mLDP Trees in GRT	661
5.17.5.2	LDP Recursive FEC Process	662
5.17.5.3	Supported Recursive Opaque Values	665
5.17.5.4	Redundancy and Resiliency	665
5.17.5.5	ASBR Physical Connection	666
5.17.5.6	OAM	666
5.17.5.7	ECMP Support.....	668
5.17.5.8	Dynamic mLDP and Static mLDP Co-existing on the Same Node.....	670
5.18	Multicast LDP Fast Upstream Switchover	671
5.18.1	Feature Configuration.....	671
5.18.2	Feature Behavior	672
5.18.3	Uniform Failover from Primary to Backup ILM.....	674
5.19	Multi-Area and Multi-Instance Extensions to LDP	675
5.19.1	LDP Shortcut for BGP Next-Hop Resolution	676
5.19.2	LDP Shortcut for IGP Routes	676
5.19.2.1	LDP Shortcut Configuration.....	677
5.19.2.2	IGP Route Resolution.....	677
5.19.2.3	LDP Shortcut Forwarding Plane	678
5.19.3	ECMP Considerations	678
5.19.4	Disabling TTL Propagation in an LSP Shortcut.....	679
5.20	LDP Graceful Handling of Resource Exhaustion.....	679
5.20.1	LDP Base Graceful Handling of Resources	680
5.21	LDP Enhanced Graceful Handling of Resources	681
5.21.1	LSR Overload Notification	682
5.21.2	LSR Overload Protection Capability	683
5.21.3	Procedures for LSR overload protection	683
5.22	User Guidelines and Troubleshooting Procedures.....	684
5.22.1	Common Procedures.....	684
5.22.2	Base Resource Handling Procedures	686
5.22.3	Enhanced Resource Handling Procedures	689
5.23	LDP IPv6 Control and Data Planes	692
5.23.1	LDP Operation in an IPv6 Network.....	693
5.23.2	Link LDP	693
5.23.3	Targeted LDP	694
5.23.4	FEC Resolution	695
5.23.5	LDP Session Capabilities	696
5.23.6	LDP Adjacency Capabilities	697
5.23.7	Address and FEC Distribution	699
5.23.8	Controlling IPv6 FEC Distribution During an Upgrade to SR OS Supporting LDP IPv6	702
5.23.9	Handling of Duplicate Link-Local IPv6 Addresses in FEC Resolution.....	703

5.23.10	IGP and Static Route Synchronization with LDP	705
5.23.11	BFD Operation.....	705
5.23.12	Services Using SDP with an LDP IPv6 FEC	706
5.23.13	Mirror Services and Lawful Intercept	707
5.23.13.1	Configuration at mirror source node	707
5.23.13.2	Configuration at mirror destination node	707
5.23.14	Static Route Resolution to a LDP IPv6 FEC.....	708
5.23.15	IGP Route Resolution to a LDP IPv6 FEC	709
5.23.16	OAM Support with LDP IPv6	709
5.23.17	LDP IPv6 Interoperability Considerations.....	710
5.23.17.1	Interoperability with Implementations Compliant with draft-ietf- mpls-ldp-ipv6	710
5.23.17.2	Interoperability with Implementations Compliant with RFC 5036 for IPv4 LDP Control Plane Only	711
5.24	LDP Process Overview.....	712
5.25	LDP-IGP Synchronization.....	713
5.26	Configuring LDP with CLI	717
5.27	LDP Configuration Overview	717
5.28	Basic LDP Configuration	717
5.29	Common Configuration Tasks	718
5.29.1	Enabling LDP.....	718
5.29.2	Configuring FEC Originate Parameters.....	719
5.29.3	Configuring Graceful-Restart Helper Parameters.....	720
5.29.4	Applying Export and Import Policies	720
5.29.5	Targeted Session Parameters.....	721
5.29.6	Interface Parameters	722
5.29.7	Session Parameters	723
5.29.8	LDP Signaling and Services	724
5.30	LDP Configuration Management Tasks.....	726
5.30.1	Disabling LDP	727
5.30.2	Modifying Targeted Session Parameters	727
5.30.3	Modifying Interface Parameters	728
5.31	LDP Command Reference	729
5.31.1	Command Hierarchies.....	729
5.31.1.1	LDP Commands	729
5.31.2	Command Descriptions	732
5.31.2.1	Generic Commands.....	732
5.31.2.2	Interface Parameters Commands.....	751
5.31.2.3	Session Parameters Commands.....	758
5.31.2.4	Targeted Session Commands	768
5.31.2.5	TCP Session Parameters Commands.....	773
5.32	Show, Clear, Debug, and Tools Command Reference	777
5.32.1	Command Hierarchies.....	777
5.32.1.1	Show Commands	777
5.32.1.2	Clear Commands.....	780
5.32.1.3	Debug Commands.....	780
5.32.1.4	Tools Commands	781
5.32.2	Command Descriptions	781
5.32.2.1	Show LDP Commands	781

5.32.2.2	Clear Commands.....	904
5.32.2.3	Debug Commands.....	906
5.32.2.4	Tools Commands	910
6	Standards and Protocol Support	917

1 Getting Started

1.1 About This Guide

This guide describes the services and protocol support provided by the router and presents examples to configure and implement MPLS, RSVP, and LDP 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.

The topics and commands described in this document apply to the:

- 7450 ESS
- 7750 SR
- 7950 XRS

[Table 1](#) lists the available chassis types for each SR OS router.

Table 1 Supported SR OS Router Chassis Types

7450 ESS	7750 SR	7950 XRS
• 7450 ESS-6/6v • 7450 ESS-7/12 running in standard mode (not mixed-mode)	• 7450 ESS-7/12 running in mixed-mode (not standard mode) • 7750 SR-a4/a8 • 7750 SR-c4/c12 • 7750 SR-1e/2e/3e • 7750 SR-7/12 • 7750 SR-12e	• 7950 XRS-16c • 7950 XRS-20/40

For a list of unsupported features by platform and chassis, refer to the *SR OS R14.0.Rx Software Release Notes*, part number 3HE10818 000x TQZZA.

Command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.



Note: This guide generically covers Release 14.0 content and may contain some content that will be released in later maintenance loads. Please refer to the *SR OS R14.0.Rx Software Release Notes*, part number 3HE10818 000x TQZZA, for information on features supported in each load of the Release 14.0 software.

1.2 Nokia Router Configuration Process

[Table 2](#) lists the tasks necessary to configure MPLS applications functions.

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

Table 2 Configuration Process

Area	Task	Chapter
Protocol configuration	Configure MPLS protocols:	
	• MPLS	MPLS
	• RSVP	RSVP
	• GMPLS	GMPLS
	• LDP	Label Distribution Protocol
Reference	List of IEEE, IETF, and other proprietary entities.	Standards and Protocol Support



Note: All features are supported on all SR OS platforms (7750 SR, 7450 ESS, and 7950 XRS) unless indicated otherwise.

2 MPLS and RSVP

2.1 In This Chapter

This chapter provides information to configure MPLS and RSVP.

- [MPLS](#)
 - [MPLS Label Stack](#)
 - [MPLS Entropy Label and Hash Label](#)
 - [Label Switching Routers](#)
 - [Bidirectional Forwarding Detection for MPLS LSPs](#)
- [RSVP](#)
 - [Using RSVP for MPLS](#)
 - [Reservation Styles](#)
 - [RSVP Overhead Refresh Reduction](#)
 - [RSVP Graceful Restart Helper](#)
 - [Enhancements to RSVP Control Plane Congestion Control](#)
 - [RSVP LSP Statistics](#)
- [MPLS Transport Profile](#)
- [Traffic Engineering](#)
 - [TE Metric \(IS-IS and OSPF\)](#)
 - [Diff-Serv Traffic Engineering](#)
 - [Diff-Serv TE LSP Class Type Change under Failure](#)
- [Advanced MPLS/RSVP Features](#)
 - [Extending RSVP LSP to use Loopback Interfaces Other than Router-id](#)
 - [LSP Path Change](#)
 - [Manual LSP Path Switch](#)
 - [Make-Before-Break \(MBB\) Procedures for LSP/Path Parameter Configuration Change](#)
 - [Automatic Creation of RSVP-TE LSP Mesh](#)
 - [RSVP-TE LSP Shortcut for IGP Resolution](#)
 - [RSVP-TE LSP Signaling using LSP Template](#)
 - [Shared Risk Link Groups](#)
 - [TE Graceful Shutdown](#)

-
- Soft Preemption of Diff-Serv RSVP LSP
 - Least-Fill Bandwidth Rule in CSPF ECMP Selection
 - Inter-Area TE LSP (ERO Expansion Method)
 - Automatic Creation of a RSVP Mesh LSP
 - Timer-based Reversion for RSVP-TE LSPs
 - Automatic Creation of an RSVP One-Hop LSP
 - MPLS Entropy Label
 - Automatic Creation of an RSVP One-Hop LSP
 - Point-to-Multipoint (P2MP) RSVP LSP
 - Application in Video Broadcast
 - P2MP LSP Data Plane
 - Ingress Path Management for P2MP LSP Packets
 - RSVP Control Plane in a P2MP LSP
 - Forwarding Multicast Packets over RSVP P2MP LSP in the Base Router
 - Segment Routing With Traffic Engineering (SR-TE)
 - SR-TE Support
 - SR-TE LSP Instantiation
 - SR-TE LSP Path Computation
 - SR-TE LSP Protection
 - Data Path Support
 - MPLS Service Usage
 - Service Distribution Paths
 - MPLS/RSVP Configuration Process Overview
 - Configuration Notes

2.2 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 not enabled by default and must be explicitly enabled.

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


2.2.1 MPLS Label Stack

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

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

As described in RFC 3032, *MPLS Label Stack Encoding*, the label stack is represented as a sequence of label stack entries. Each label stack entry is represented by 4 octets. [Figure 1](#) displays the label placement in a packet.

Figure 1 Label Placement




OSSG013

Table 3 Packet/Label Field Description

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

A stack can carry several labels, organized in a last in/first out order. The top of the label stack appears first in the packet and the bottom of the stack appears last, as shown in [Figure 2](#).

Figure 2 Label Packet Placement



OSSG014

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

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

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

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

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

2.2.1.1 Label Values

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

- A value of 0 represents the IPv4 Explicit NULL Label. This Label value is legal only at the bottom of the Label stack. It indicates that the Label stack must be popped, and the packet forwarding must be based on the IPv4 header.

- A value of 1 represents the router alert Label. This Label value is legal anywhere in the Label stack except at the bottom. When a received packet contains this Label value at the top of the Label stack, it is delivered to a local software module for processing. The actual packet forwarding is determined by the Label beneath it in the stack. However, if the packet is further forwarded, the router alert Label should be pushed back onto the Label stack before forwarding. The use of this Label is analogous to the use of the router alert option in IP packets. Since this Label cannot occur at the bottom of the stack, it is not associated with a particular network layer protocol.
- A value of 2 represents the IPv6 explicit NULL Label. This Label value is only legal at the bottom of the Label stack. It indicates that the Label stack must be popped, and the packet forwarding must be based on the IPv6 header.
- A value of 3 represents the Implicit NULL Label. This is a Label that a Label Switching Router (LSR) can assign and distribute, but which never actually appears in the encapsulation. When an LSR would otherwise replace the Label at the top of the stack with a new Label, but the new Label is Implicit NULL, the LSR pops the stack instead of doing the replacement. Although this value may never appear in the encapsulation, it needs to be specified in the Label Distribution Protocol (LDP), so a value is reserved.
- Values 4 to 15 are reserved for future use.

The router uses labels for MPLS, RSVP-TE, and LDP, as well as packet-based services such as VLL and VPLS.

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

- Label values 16 through 31 are reserved for future use.
- Label values 32 through 1,023 are available for static LSP label assignments.
- Label values 1,024 through 2,047 are reserved for future use.
- Label values 2,048 through 18,431 are available for static service label assignments
- Label values 18,432 through 262,143 (131,071 in chassis modes lower than D) are assigned dynamically by RSVP, LDP, and BGP control planes for both MPLS LSP and service labels.
- Label values 262,144 (131,072 in chassis modes lower than D) through 1,048,575 are reserved for future use.

2.2.2 MPLS Entropy Label and Hash Label

The router supports both the MPLS entropy label (RFC 6790) and the flow-aware transport label (known as the hash label) (RFC 6391). These labels allows LSR nodes in a network to load-balance labeled packets in a much more granular fashion than allowed by simply hashing on the standard label stack. The labels also remove the need to have an LSR inspect the payload below the label stack to check for an IPv4 or IPv6 header.

The application of a hash label or an entropy label is mutually exclusive for a service.

The hash label is supported on VLL, VPRN, or VPLS services bound to any MPLS type encapsulated SDPs as well as to a VPRN service using **auto-bind-tunnel** with the **resolution-filter** set to any MPLS tunnel type. When this feature is enabled, the ingress data path is modified such that the result of the hash on the payload packet header is communicated to the egress data path for use as the value of the label field of the hash label. The egress data path appends the hash label to the bottom of the stack (BoS) and sets the S-bit to 1. The user enables the signaling of the hash-label capability under a VLL spoke sdp, a VPLS spoke sdp or mesh sdp, or an IES/VPRN spoke SDP interface by adding the **signal-capability** option. When this capability is enabled, the decision to insert the hash label on the user and control plane packets by the local PE is solely determined by the outcome of the signaling process and may override the local PE configuration.

Entropy labels are supported for the following services and protocols:

- EVPN VPLS and Epipe
- RFC 3107 MP-BGP tunnels
- VLLs (including BGP VPWS), IES/VPRN and VPLS spoke SDP termination, but not including Apipes and Cpipes
- LDP VPLS and BGP-AD VPLS

The MPLS entropy label provides a similar function to the hash label but is applicable to a wider range of services. The entropy label is appended directly below the tunnel label (which may be above the bottom of the stack). As with the hash label, the value of the entropy label is calculated based on a hash of the packet payload header.

The entropy label is inserted along with an entropy label indicator (ELI). The ELI is a special-purpose MPLS label (Value = 7) that indicates that the entropy label follows in the stack. It is always placed immediately below the tunnel label to which hashing applies.

The ability of a node to receive and process an entropy label for an LSP tunnel is signaled using capability signaling. Entropy labels are supported on RSVP, LDP, and BGP tunnels. The entropy label is not supported on P2MP LSPs. Entropy label capability is advertised at the tunnel level by the downstream LSR. This indicates the capability of the node to receive and process the entropy label. This capability can be advertised for an LDP, RSVP, or BGP FEC. Capability signaling is not supported for BGP tunnels, as no agreed standard for this exists in the IETF. The router therefore supports manual configuration of an override of entropy label capability for a BGP tunnel at the ingress LER. An LSR for RSVP and LDP tunnels will pass the entropy label capability from the downstream LSP segment to upstream peers. However, earlier releases that do not support entropy label functionality will not pass the entropy label capability to their peers.

The insertion of an entropy label by the upstream LER on a tunnel for which entropy label capability is known is configured on a per-service or per-protocol basis. It is only inserted if the downstream peer actually signals entropy label support (or the entropy label capability is overridden in the case of a BGP ingress LER). The router will only insert a single entropy label even if multiple LSP labels exist in a label stack. The entropy label and its ELI are inserted immediately below the innermost tunnel label for which the downstream peer LER is known to be entropy-label-capable. This ensures that the entropy label is preserved as far as possible along a path through the network.

The router supports the Entropy Label for the following services/protocols:

- VPRN
- EVPN VPLS and Epipe
- RFC 3107 MP-BGP tunnels
- RSVP and LDP LSPs used as shortcuts for static, IGP and BGP route resolution
- VLLs, including BGP VPWS, IES/VPRN and VPLS spoke-sdp termination, but not including Apipe and Cpipe
- LDP VPLS and BGP-AD VPLS

Entropy label insertion is configured using the **entropy-label** command under the service or spoke SDP context. This command and the **hash-label** command are mutually exclusive.

The entropy label and ELI are not supported on pseudowire ports. That is, although the LSP used by a pseudowire port may signal entropy label capability (because other services may use the same LSP), the entropy label/ELI is not inserted on packets from a pseudowire SAP using a pseudowire port.

If entropy label insertion is configured for a VPLS or VLL service, the MTU of the SDP binding is automatically reduced to account for the overhead of the two labels (entropy label and ELI). This will happen whether or not the LSP tunnel used by the service is entropy-label-capable.

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 a user to avoid entropy label insertion where there is a risk of the label stack depth becoming too great.

2.2.2.1 Inserting and Processing the Entropy Label

2.2.2.1.1 Ingress LER

The procedures at the ingress LER are as specified in Section 4.2 of RFC 6790. In general, the router inserts an entropy label in a packet if the downstream node for the LSP tunnel has signaled support for entropy labels, the entropy label is configured for the service that the packet belongs to, and the entropy label is not disabled for an RSVP LSP.

RFC 6790 suggests that the ingress LER may 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), the router only inserts a single entropy label/ELI pair under the innermost LSP label closest to the service payload that has advertised entropy label capability.

The router does not insert an entropy label in a packet belonging to a service for which the hash label has been configured, even if the far end for the LSP tunnel has advertised entropy label capability. The system will instead insert a hash label, as specified by the hash label feature.

The router inserts an entropy label on a tunnel that is ELC (if the far-end node has signaled entropy label capability for LDP or RSVP, or if the entropy label capability override is configured for BGP tunnels) when required by the service (as determined by the **entropy-label** command), even if an implicit or explicit NULL label has been signaled by the downstream LSR or LER. This ensures consistent behavior as well as ensuring that entropy as determined by the ingress LER is maintained where a tunnel with an implicit NULL label is stitched at a downstream LSR.

2.2.2.1.2 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**: Only use the entropy label as input to the hash routine. Ignore the rest of the label stack.
- **label-ip**: Only use the entropy label and the IP packet as input to the hash routine. Ignore the rest of the label stack.

The presence of an entropy label is ignored if IP-only load balancing is configured.

If 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.

2.2.2.2 Egress LER

If an entropy label is detected in the label stack at an egress LER for a tunnel where the tunnel label that the entropy label is associated with is popped, then the entropy label is also popped and the packet processed as normal. This occurs whether or not the system has signalled entropy-label-capability.

If an ELI is popped that has the BoS bit set, then the system should discard the packet and raise a trap as per Section 4.1 of RFC 6790.

2.2.2.3 Mapping Entropy Label and Entropy Label Capability at LSP Stitching Points

An entropy label and ELI will not be exposed when a tunnel label is swapped at an LSR acting as an LSP stitching point. The entropy label and ELI will therefore be forwarded as any other packet on the LSP.

A router acting as a stitching point between two LSPs will map the entropy label capability received in signaling for a downstream segment to the upstream segment for the level in the LSP hierarchy being stitched.

If **override-tunnel-elc** is configured for a downstream BGP segment, then the system will signal entropy label capability upstream on all RSVP LSP segments that are stitched to a downstream BGP tunnel segment. This must be configured correctly; otherwise packets with an entropy label will be discarded by a downstream LER that is not entropy-label-capable.

The mapping of entropy label capability across LDP-BGP stitching points is not supported. If a downstream tunnel endpoint signals entropy label capability, this will not be propagated to the ingress LER. The entropy label/ELI should not be inserted on these LSPs by the ingress LER.

2.2.2.4 Entropy Label on OAM Packets

Service OAM packets or within the context of a shortcut, for example, ICMP Ping/traceroute packet, also include an entropy label/ELI if ELC is signaled for the corresponding tunnel and **entropy-label** is enabled for the service. The entropy label/ELI is inserted at the same level in the label stack as it is in user data packets; which is under the innermost LSP label closest to the service payload that has advertised entropy label capability. The entropy label/ELI 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 entropy label/ELI inserted.

2.2.3 Label Switching Routers

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

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

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

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

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

2.2.3.1 LSP Types

The following are LSP types:

- Static LSPs — A static LSP specifies a static path. All routers that the LSP traverses must be configured manually with labels. No signaling such as RSVP or LDP is required.
- Signaled LSP — LSPs are set up using a signaling protocol such as RSVP-TE or LDP. The signaling protocol allows labels to be assigned from an ingress router to the egress router. Signaling is triggered by the ingress routers. Configuration is required only on the ingress router and is not required on intermediate routers. Signaling also facilitates path selection.

There are two signaled LSP types:

- Explicit-path LSPs — MPLS uses RSVP-TE to set up explicit path LSPs. The hops within the LSP are configured manually. The intermediate hops must be configured as either strict or loose meaning that the LSP must take either a direct path from the previous hop router to this router (strict) or can traverse through other routers (loose). You can control how the path is set up. They are similar to static LSPs but require less configuration. See [RSVP](#).
- Constrained-path LSPs — The intermediate hops of the LSP are dynamically assigned. A constrained path LSP relies on the Constrained Shortest Path First (CSPF) routing algorithm to find a path which satisfies the constraints for the LSP. In turn, CSPF relies on the topology database provided by the extended IGP such as OSPF or IS-IS.
Once the path is found by CSPF, RSVP uses the path to request the LSP set up. CSPF calculates the shortest path based on the constraints provided such as bandwidth, class of service, and specified hops.

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

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

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

2.2.4 Bidirectional Forwarding Detection for MPLS LSPs

BFD for MPLS LSPs enables BFD to monitor the LSP between its LERs, irrespective of how many LSRs the LSP may traverse. Therefore, it enables faults that are local to individual LSPs to be detected, whether or not they also affect forwarding for other LSPs or IP packet flows. This makes it ideal for monitoring LSPs carrying specific high-value services, where detecting forwarding failures in the minimal amount of time is critical. The system will raise an SNMP trap, as well as indicate the BFD session state in show and tools dump commands if a LSP BFD session goes down.

The system supports LSP BFD on RSVP LSPs. BFD packets are encapsulated in an MPLS label stack corresponding to the FEC that the BFD session is associated with, as described in Section 7 of RFC5884.

Since RSVP LSPs are unidirectional, a routed return path is used for the BFD control packets from the egress LER towards the ingress LER.

2.2.4.1 Bootstrapping the BFD Session

A BFD session on an LSP is bootstrapped using LSP Ping. LSP Ping is used to exchange the local/remote discriminator values to use for the BFD session for a particular <MPLS LSP, FEC>.

The process is as follows:

1. The ingress LER sends an LSP ping to egress LSR, including the BFD Discriminator TLV with the local discriminator value.
2. The egress LER validates the received FEC. If this is true, then the Egress LSR sends a BFD control packet to the ingress LSR with the your discriminator field set to the value of the local discriminator in the received LSP Ping, and my discriminator field set to the locally assigned discriminator value. This BFD control packet will be routed to the ingress LSR.
3. The ingress LER then follows normal BFD procedures in response to the BFD control packets from the egress LSR.

The LSP ping interval is configured using the **lsp-ping-interval** command under the **bfd** context for the LSP or LSP primary path.

2.2.4.2 LSP BFD Configuration

There are four steps to configuring LSP BFD:

1. Configure BFD Template
2. Enable LSP BFD on the tail node or configure the maximum number of LSP BFD sessions at the tail node.
3. Apply BFD Template to the LSP or LSP Path.
4. Enable BFD on the LSP or LSP Path.

LSP BFD uses BFD templates to set generic BFD session parameters.

The BFD template is configured as follows:

```
config
router
  bfd
    bfd-template name
      transmit-interval transmit-interval
      receive-interval receive-interval
      echo-receive echo-interval
      multiplier multiplier
      type cpm-np
    exit
```

Network processor BFD is not supported for LSPs. Furthermore, the minimum supported receive or transmit timer interval is 1 second. Therefore, an error will be generated if a user tries to bind a **bfd-template** with the '**type cpm-np**' command or any unsupported transmit or receive interval value to an LSP. An error will also be generated as user attempts to commit changes to a BFD template that is already bound to an LSP where the new values are invalid for `lsp-bfd`.

BFD templates may be used by different BFD applications (for example, LSPs or pseudowires). If the BFD timer values are changed in a template, the BFD sessions on LSPs or spoke-SDPs to which that template is bound will try to renegotiate their timers to the new values.

The `bfd-template` uses a begin-commit model. To edit any value within the BFD template, a `<begin>` needs to be executed before the template context has been entered. However, a value will still be stored temporarily in the template-module until the commit is issued. Values will actually be used once the commit is issued.

2.2.4.3 Enabling and Implementing Limits for LSP BFD on a Node

The **`config>router>lsp-bfd`** command enables support for LSP BFD and allows an upper limit to the number of supported sessions at the tail end node for LSPs, where it is disabled by default. This is useful because BFD resources are shared among applications using BFD, so a user may wish to set an upper limit to ensure that a certain number of BFD sessions are reserved for other applications. This is important at the tail end of LSPs where no per-LSP configuration context exists.

LSP BFD is enabled or disabled on a node-wide basis using the **`bfd-sessions max-limit`** command under the **`config>router>lsp-bfd`** context. This command also enables the maximum number of LSP BFD sessions that can be established at the tail end of LSPs to be limited.

The default is disabled. The *max-limit* parameter specifies the maximum number of LSP BFD sessions that the system will allow to be established at the tail end of LSPs.

2.2.4.4 BFD Configuration on RSVP-TE LSPs

LSP BFD is applicable to configured RSVP LSPs as well as mesh-p2p and one-hop-p2p auto-LSPs.

It is configured on an RSVP-TE LSP, or on the primary path of an RSVP-TE LSP, under a new **bfd** context at the LSP head end.

A BFD template must always be configured first. BFD is then enabled using the **bfd-enable** command.

```
config
router
mpls
lsp xyz
bfd
[no] bfd-template name
[no] bfd-enable
exit
```

When BFD is configured at the LSP level, BFD packets follow the currently active path of the LSP.

The **bfd-template** provides the control packet timer values for the BFD session to use at the LSP head end. Since there is no LSP configuration at the tail end of an RSVP LSP, the BFD state machine at the tail end will initially use system-wide default parameters (the timer values are: min-tx: 1sec, min-rx: 1sec). The head end will then attempt to adjust the control packet timer values when it transitions to the INIT state.

BFD is configured at the primary path level, as follows:

```
config
router
mpls
lsp <xyz>
primary path-name
bfd
[no] bfd-template name
[no] bfd-enable
exit
```

It is not possible to configure LSP BFD on a secondary path or on P2MP LSPs.

LSP BFD at the LSP level and the path level is mutually exclusive. That is, if LSP BFD is already configured for the LSP then its configuration for the path is blocked. Likewise it cannot be configured on the LSP if it is already configured at the path level.

LSP BFD is supported on auto-LSPs. In this case, LSP BFD is configured on mesh-p2p and one-hop-p2p auto-LSPs using the LSP template, as follows:

```
Config
router
mpls
  lsp-template template-name {mesh-p2p | one-hop-p2p}
  bfd
    [no] bfd-template name
    [no] bfd-enable
  exit
```

2.2.4.5 MPLS Facility Bypass Method of MPLS Fast Re-Route (FRR)

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

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

- When a down stream detour becomes active at a point of local repair (PLR):
The ingress LER switches to the standby LSP path. If the primary LSP path is repaired subsequently at the PLR, the LSP will switch back to the primary path. If the standby goes down, the LSP is switched back to the primary, even though it is still on the detour at the PLR. If the primary goes down at the ingress while the LSP is on the standby, the detour at the ingress is cleaned up and for one-to-one detours a “path tear” is sent for the detour path. In other words, the detour at the ingress does not protect the standby. If and when the primary LSP is again successfully re-signaled, the ingress detour state machine will be restarted.
- When the primary fails at the ingress:
The LSP switches to the detour path. If a standby is available then LSP would switch to standby on expiration of **hold-timer**. If **hold-timer** is disabled then switchover to standby would happen immediately. On successful global revert of primary path, the LSP would switch back to the primary path.
- Admin groups are not taken into account when creating detours for LSPs.

2.2.4.6 Manual Bypass LSP

In prior releases, the router implemented dynamic bypass tunnels as per RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*. When an LSP is signaled and the local protection flag is set in the session_attribute object and/or the FRR object in the path message indicates that facility backup is desired, the PLR will establish a bypass tunnel to provide node and link protection. If a bypass LSP which merges in a downstream node with the protected LSP exist, and if this LSP satisfies the constraints in the FRR object, then this bypass tunnel is selected.

With the manual bypass feature, an LSP can be preconfigured from a PLR which will be used exclusively for bypass protection. When a path message for a new LSP requests bypass protection, the node will first check if a manual bypass tunnel satisfying the path constraints exists. If one is found, it will be selected. If no manual bypass tunnel is found, the router will dynamically signal a bypass LSP in the default behavior. Users can disable the dynamic bypass creation on a per node basis using the CLI.


A maximum of 1000 associations of primary LSP paths can be made with a single manual bypass by default. The max-bypass-associations integer command increases the number of associations. If dynamic bypass creation is disabled on the node, it is recommended to configure additional manual bypass LSPs to handle the required number of associations.

Refer to [Configuring Manual Bypass Tunnels](#) for configuration information.

2.2.4.6.1 PLR Bypass LSP Selection Rules

The PLR uses rules to select a bypass LSP among multiple manual and dynamic bypass LSPs at the time of establishment of the primary LSP path or when searching for a bypass for a protected LSP which does not have an association with a bypass tunnel: [Figure 3](#) shows an example of bypass tunnel nodes.

Figure 3 Bypass Tunnel Nodes



al_0204

The rules are:

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

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

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


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

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

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

2.2.4.6.2 FRR Node-Protection (Facility)

The MPLS Fast Re-Route (FRR) functionality enables PLRs to be aware of the missing node protection and lets them regularly probe for a node-bypass. [Figure 4](#) shows an example of FRR node protection.

Figure 4 FRR Node-Protection Example

The following describes an LSP scenario where:

- LSP 1: between PE_1 to PE_2, with CSPF, FRR facility node-protect enabled.
- P_1 protects P_2 with bypass-nodes P_1 - P_3 - P_4 - PE_4 - PE_3.
- If P_4 fails, P_1 tries to establish the bypass-node three times.
- When the bypass-node creation fails, P_1 will protect link P_1-P_2.
- P_1 protects the link to P_2 through P_1 - P_5 - P_2.
- P_4 returns online.

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

2.2.4.7 Uniform FRR Failover Time

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

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

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

2.2.4.8 MPLS/RSVP on Broadcast Interface

The MPLS/RSVP on Broadcast Interface feature allows MPLS/RSVP to distinguish neighbors from one another when the outgoing interface is a broadcast interface connecting to multiple neighbors over a broadcast domain. More specifically, in the case where a BFD session towards a specific neighbor on the broadcast domain goes down, the consecutive actions (for example, FRR switchover) will only concern the LSPs of the affected neighbor. Previously, the actions would have been taken on the LSPs of all neighbors over the outgoing interface.

2.2.4.9 Automatic Bandwidth Allocation for RSVP LSPs

This section includes the following topics:

- [Enabling and Disabling Auto-Bandwidth Allocation on an LSP](#)
- [Autobandwidth on LSPs with Secondary or Secondary Standby Paths](#)
- [Measurement of LSP Bandwidth](#)
- [Passive Monitoring of LSP Bandwidth](#)
- [Periodic Automatic Bandwidth Adjustment](#)
- [Overflow-Triggered Auto-Bandwidth Adjustment](#)
- [Manually-Triggered Auto-Bandwidth Adjustment](#)

2.2.4.9.1 Enabling and Disabling Auto-Bandwidth Allocation on an LSP

This section discusses an auto-bandwidth hierarchy configurable in the **config>router>mpls>lsp** context.

Adding auto-bandwidth at the LSP level starts the measurement of LSP bandwidth described in [Measurement of LSP Bandwidth](#) and allows auto-bandwidth adjustments to take place based on the triggers described in [Periodic Automatic Bandwidth Adjustment](#).

When an LSP is first established, the bandwidth reserved along its primary path is controlled by the bandwidth parameter in the **config>router>mpls>lsp>primary** context, whether or not the LSP has auto-bandwidth enabled, while the bandwidth reserved along a secondary path is controlled by the bandwidth parameter in the **config>router>mpls>lsp>secondary** context. When auto-bandwidth is enabled and a trigger occurs, the system will attempt to change the bandwidth of the LSP to a value between min-bandwidth and max-bandwidth, which are configurable values in the **lsp>auto-bandwidth** context. **min-bandwidth is the minimum bandwidth that auto-bandwidth can signal for the LSP and max-bandwidth is the maximum bandwidth that can be signaled. The user can set the min-bandwidth to the same value as the primary path bandwidth but the system will not enforce this restriction. The system will allow:**

- No **min-bandwidth** to be configured. In this case, the implicit minimum is 0 Mb/s
- No **max-bandwidth** to be configured, as long as overflow-triggered auto-bandwidth is not configured. In this case, the implicit maximum is infinite (effectively 100 Gbps).
- The configured primary path bandwidth to be outside the range of min-bandwidth to max-bandwidth.
- **auto-bandwidth** parameters can be changed at any time on an operational LSP; in most cases the changes have no immediate impact but subsequent sections will describe some exceptions

All of the auto-bandwidth adjustments discussed are performed using MBB procedures.

Auto bandwidth can be added to an operational LSP at any time (without the need to shut down the LSP or path), but no bandwidth change occurs until a future trigger event. Auto bandwidth may also be removed from an operational LSP at any time and this causes an immediate MBB bandwidth change to be attempted using the configured primary path bandwidth.

A change to the configured bandwidth of an auto-bandwidth LSP has no immediate affect. The change only occurs if the LSP/path goes down (due to failure or administrative action) and comes back up, or if auto-bandwidth is removed from the LSP. The operator can force an auto-bandwidth LSP to be resized immediately to an arbitrary bandwidth using the appropriate tools commands.

2.2.4.9.2 Autobandwidth on LSPs with Secondary or Secondary Standby Paths

Autobandwidth is supported for LSPs that have secondary or secondary standby paths. A secondary path is only initialized at its configured bandwidth when it is established, and the bandwidth is adjusted only when the secondary path becomes active.

This description makes use of the following terminology:

- **current_BW:** the last known reserved bandwidth for the LSP; may be the value of a different path from the currently active path
- **operational BW:** the last known reserved BW for a given path, as recorded in the MIB
- **configured BW:** the bandwidth explicitly configured for the LSP path by the user in CLI
- **active path:** the path (primary or secondary) the LSP currently uses to forward traffic
- **signaled BW:** the new BW value signaled during an MBB

A secondary or standby secondary path is initially signaled with its configured bandwidth. Setup for the secondary path is triggered only when the active path goes down or becomes degraded (e.g. due to FRR or preemption). An auto-BW triggered bandwidth adjustment (auto bandwidth MBB) only takes place on the active path. For example, if an auto-BW adjustment occurs on the primary path, which is currently active, no adjustment is made at that time to the secondary path since that path is not active.

When the active path changes, the `current_bw` is updated to the operational bandwidth of the newly active path. While the auto-BW MBB on the active path is in progress, a statistics sample could be triggered, and this would be collected in the background. Auto-bandwidth computations will use the `current_bw` of the newly active path. In case the statistics sample collection results in a bandwidth adjustment, the in-progress auto-BW MBB will be restarted. If after five attempts, the auto-BW MBB fails, the `current_bw` and secondary operational BW remain unchanged.

For a secondary or standby secondary path, if the active path for an LSP changes (without the LSP going down), an auto-BW MBB will be triggered for the new active path. The bandwidth used to signal the MBB will be the operational bandwidth of the previous active path. If the MBB fails, it will retry with a maximum of 5 attempts. The reserved bandwidth of the newly active path will therefore be its configured bandwidth until the MBB succeeds.

For a secondary path where the active path goes down, the LSP will go down temporarily until the secondary path is setup. If the LSP goes down, all statistics and counters are cleared, so the previous path operational bandwidth is lost. That is, the operational BW of a path is not persistent across LSP down events. In this case, there will be no immediate bandwidth adjustment on the secondary path.

The following algorithm is used to determine the signaled bandwidth on a newly active path:

1. For a path that is operationally down, signaled_bw = config_bw.
2. For the active path, if an auto-BW MBB adjustment is in progress, signaled_bw = previous path operational BW for the first 5 attempts. For the remaining attempts, the signaled BW = operational BW.
3. For an MBB on the active path (other than an auto-BW MBB), MBB signaled BW = operational BW.
4. For an MBB on the inactive path, MBB signaled BW = configured BW.

If the primary path is not the currently active path and it has not gone down, then any MB uses the configured BW for the primary path. However, if the configured BW is changed for a path that is currently not active, then a config change MBB is not triggered.

If the standby is SRLG enabled, and the active path is the standby, and the primary comes up, this will immediately trigger a delayed retry MBB on the standby. If the delayed retry MBB fails, immediate reversion to the primary will occur irrespective of the retry timer.

When the system reverts from a secondary standby or secondary path to the primary path, a Delayed Retry MBB will be attempted to bring bandwidth of the standby path back to its configured bandwidth. Delayed Retry MBB is attempted once, and if it fails, the standby will be torn down. A Delayed Retry MBB has highest priority among all MBBs, so it will take precedence over any other MBB in progress on the standby path (e.g. Config change or Preemption).

The system will carry-over the last signaled BW of the LSP over multiple failovers. For example, if an LSP is configured with auto-BW for some time, and adjusts its currently reserved bandwidth for the primary, and Monitor mode is then enabled, BW adjustment on the primary ceases, but the BW remains reserved at the last adjusted value. Next, the LSP fails over to a secondary or secondary standby. The secondary will inherit the last reserved BW of the primary, but then disable further adjustment as long as monitoring mode is enabled.

The system's ability to carry-over the last signaled BW across failovers has the following limitations:

- Case 1: If the LSP fails over from path1 to path2 and the AutoBW MBB on path2 is successful, the last signaled BW is carried over when the LSP reverts back to path1 or fails over to a new path3. This may trigger an AutoBW MBB on the new active path to adjust its bandwidth to last signaled BW.
- Case 2: If the LSP fails over from path1 to path2 and the AutoBW MBB on path2 is still in progress and the LSP reverts back to path1 or fails over to a new path3, the last signaled BW is carried over to the new active path (path1 or path3) and this may result in an AutoBW MBB on that path.
- Case 3: If the LSP fails over from path1 to path2 and the AutoBW MBB on path2 fails (after 5 retry attempts), the last signaled BW from when path1 was active is lost. Therefore, when the LSP reverts back to path1 or fails over to a new path3, the original signaled BW from path1 is not carried over. However the signaled bandwidth of path2 will be carried over to the new active path (path1 or path3) and may trigger an AutoBW on that path.

2.2.4.9.3 Measurement of LSP Bandwidth

Automatic adjustment of RSVP LSP bandwidth based on measured traffic rate into the tunnel requires the LSP to be configured for egress statistics collection at the ingress LER. The following CLI shows an example:

```
config router mpls lsp name
    egress-statistics
    accounting-policy 99
    collect-stats
    no shutdown
exit
```

All LSPs configured for accounting, including any configured for auto-bandwidth based on traffic measurements, must reference the same accounting policy. An example configuration of such an accounting-policy is shown below: in the CLI example below.

```
config log
    accounting-policy 99
```

```

collection-interval 5
record combined-mpls-lsp-egress
exit
exit

```

The record **combined-mpls-lsp-egress** command in the accounting policy has the effect of recording both egress packet and byte counts and bandwidth measurements based on the byte counts if auto-bandwidth is enabled on the LSP.

When egress statistics are enabled the CPM collects stats from of all XCMs or IOMs involved in forwarding traffic belonging to the LSP (whether the traffic is currently leaving the ingress LER via the primary LSP path, a secondary LSP path, an FRR detour path or an FRR bypass path). The egress statistics have counts for the number of packets and bytes forwarded per LSP on a per-forwarding class, per-priority (in-profile vs. out-of-profile) basis. When auto-bandwidth is configured for an LSP the ingress LER calculates a traffic rate for the LSP as follows:

Average data rate of LSP[x] during interval[i] = $F(x, i) - F(x, i-1) / \text{sample interval}$

$F(x, i)$ — The total number of bytes belonging to LSP[x], regardless of forwarding-class or priority, at time[i]

sample interval = time[i] — time [i-1], time[i+1] — time[i], etc.

The sample interval is the product of sample-multiplier and the collection-interval specified in the auto-bandwidth accounting policy. A default sample-multiplier for all LSPs may be configured using the **config>router>mpls>auto-bandwidth-defaults** command but this value can be overridden on a per-LSP basis at the **config>router>mpls>lsp>auto-bandwidth** context. The default value of sample-multiplier (the value that would result from the no auto-bandwidth-defaults command) is 1, which means the default sample interval is 300 seconds.

Over a longer period of time called the adjust interval the router keeps track of the maximum average data rate recorded during any constituent sample interval. The adjust interval is the product of adjust-multiplier and the collection-interval specified in the auto-bandwidth accounting-policy. A default adjust-multiplier for all LSPs may be configured using the **config>router>mpls>auto-bandwidth-multiplier** command but this value can be overridden on a per-LSP basis at the **config>router>mpls>lsp>auto-bandwidth** context. The default value of adjust-multiplier (the value that would result from the no auto-bandwidth-multitplier command) is 288, which means the default adjust interval is 86400 seconds or 24 hours. The system enforces the restriction that adjust-multiplier is equal to or greater than sample-multiplier. It is recommended that the adjust-multiplier be an integer multiple of the sample-multiplier.

The collection-interval in the auto-bandwidth accounting policy can be changed at any time, without disabling any of the LSPs that rely on that policy for statistics collection.

The sample-multiplier (at the **mpls>auto-bandwidth** level or the **lsp>auto-bandwidth** level) can be changed at any time. This will have no effect until the beginning of the next sample interval. In this case the adjust-interval does not change and information about the current adjust interval (such as the remaining adjust-multiplier, the maximum average data rate) is not lost when the sample-multiplier change takes effect.

The system allows adjust-multiplier (at the **mpls** level or the **lsp>auto-bandwidth** level) to be changed at any time as well but in this case the new value shall have no effect until the beginning of the next adjust interval.

Byte counts collected for LSP statistics include layer 2 encapsulation (Ethernet headers and trailers) and therefore average data rates measured by this feature include Layer 2 overhead as well.

2.2.4.9.4 Passive Monitoring of LSP Bandwidth

The system offers the option to measure the bandwidth of an RSVP LSP (see [Measurement of LSP Bandwidth](#)) without taking any action to adjust the bandwidth reservation, regardless of how different the measured bandwidth is from the current reservation. Passive monitoring is enabled using the **config>router>mpls>lsp>auto-bandwidth>monitor-bandwidth** command.

The **show>router>mpls>lsp detail** command can be used to view the maximum average data rate in the current adjust interval and the remaining time in the current adjust interval.

2.2.4.9.5 Periodic Automatic Bandwidth Adjustment

Automatic bandwidth allocation is supported on any RSVP LSP that has MBB enabled. MBB is enabled in the **config>router>mpls>lsp** context using the **adaptive** command. For automatic adjustments of LSP bandwidth to occur the **monitor-bandwidth** command must not be present at **config>router>mpls>lsp>auto-bandwidth** context, otherwise only passive measurements will occur.

If an eligible RSVP LSP is configured for auto-bandwidth, by entering auto-bandwidth at the config>router>mpls>lsp context, then the ingress LER decides every adjust interval whether to attempt auto-bandwidth adjustment. The following parameters are defined:

- `current_bw` — The currently reserved bandwidth of the LSP; this is the operational bandwidth that is already maintained in the MIB.
- `measured_bw` — The maximum average data rate in the current adjust interval.
- `signaled_bw` — The bandwidth that is provided to the CSPF algorithm and signaled in the SENDER_TSPEC and FLOWSPEC objects when an auto-bandwidth adjustment is attempted.
- `min` — The configured min-bandwidth of the LSP.
- `max` — The configured max-bandwidth of the LSP.
- `up%` — The minimum difference between `measured_bw` and `current_bw`, expressed as a percentage of `current_bw`, for increasing the bandwidth of the LSP.
- `up` — The minimum difference between `measured_bw` and `current_bw`, expressed as an absolute bandwidth relative to `current_bw`, for increasing the bandwidth of the LSP. This is an optional parameter; if not defined the value is 0.
- `down%` — The minimum difference between `current_bw` and `measured_bw`, expressed as a percentage of `current_bw`, for decreasing the bandwidth of the LSP.
- `down` — The minimum difference between `current_bw` and `measured_bw`, expressed as an absolute bandwidth relative to `current_bw`, for decreasing the bandwidth of the LSP. This is an optional parameter; if not defined the value is 0.

At the end of every adjust interval the system decides if an auto-bandwidth adjustment should be attempted. The heuristics are as follows:

- If the measured bandwidth exceeds the current bandwidth by more than the percentage threshold and also by more than the absolute threshold then the bandwidth is re-signaled to the measured bandwidth (subject to min and max constraints).
- If the measured bandwidth is less than the current bandwidth by more than the percentage threshold and also by more than the absolute threshold then the bandwidth is re-signaled to the measured bandwidth (subject to min and max constraints).
- If the current bandwidth is greater than the max bandwidth then the LSP bandwidth is re-signaled to max bandwidth, even if the thresholds have not been triggered.

- If the current bandwidth is greater than the min bandwidth then the LSP bandwidth is re-signaled to min bandwidth, even if the thresholds have not been triggered.

Changes to min-bandwidth, max-bandwidth and any of the threshold values (up, up%, down, down%) are permitted at any time on an operational LSP but the changes have no effect until the next auto-bandwidth trigger (for example, adjust interval expiry).

If the measured bandwidth exceeds the current bandwidth by more than the percentage threshold and also by more than the absolute threshold then the bandwidth is re-signaled to the measured bandwidth (subject to min and max constraints).

The adjust-interval and maximum average data rate are reset whether the adjustment succeeds or fails. If the bandwidth adjustment fails (for example, CSPF cannot find a path) then the existing LSP is maintained with its existing bandwidth reservation. The system does not retry the bandwidth adjustment (for example, per the configuration of the LSP retry-timer and retry-limit).

2.2.4.9.6 Overflow-Triggered Auto-Bandwidth Adjustment

For cases where the measured bandwidth of an LSP has increased significantly since the start of the current adjust interval it may be desirable for the system to preemptively adjust the bandwidth of the LSP and not wait until the end of the adjust interval.

The following parameters are defined:

- `current_bw` — The currently reserved bandwidth of the LSP.
- `sampled_bw` — The average data rate of the sample interval that just ended.
- `measured_bw` — The maximum average data rate in the current adjust interval.
- `signaled_bw` — The bandwidth that is provided to the CSPF algorithm and signaled in the `SENDER_TSPEC` and `FLOWSPEC` objects when an auto-bandwidth adjustment is attempted.
- `max` — The configured max-bandwidth of the LSP.
- `%_threshold` — The minimum difference between `sampled_bw` and `current_bw`, expressed as a percentage of the `current_bw`, for counting an overflow event.
- `min_threshold` — The minimum difference between `sampled_bw` and `current_bw`, expressed as an absolute bandwidth relative to `current_bw`, for counting an overflow event. This is an optional parameter; if not defined the value is 0.

When a sample interval ends it is counted as an overflow if:

- The sampled bandwidth exceeds the current bandwidth by more than the percentage threshold and by more than the absolute bandwidth threshold (if defined).
- When the number of overflow samples reaches a configured limit, an immediate attempt is made to adjust the bandwidth to the measured bandwidth (subject to the min and max constraints).

If the bandwidth adjustment is successful then the adjust-interval, maximum average data rate and overflow count are all reset. If the bandwidth adjustment fails then the overflow count is reset but the adjust-interval and maximum average data rate continue with current values. It is possible that the overflow count will once again reach the configured limit before the end of adjust-interval is reached and this will once again trigger an immediate auto-bandwidth adjustment attempt.

The overflow configuration command fails if the max-bandwidth of the LSP has not been defined.

The threshold limit can be changed on an operational auto-bandwidth LSP at any time and the change should take effect at the end of the current sample interval (for example, if the user decreases the overflow limit to a value lower than the current overflow count then auto-bandwidth adjustment will take place as soon as the sample interval ends). The threshold values can also be changed at any time (for example, %_threshold and min_threshold) but the new values will not take effect until the end of the current sample interval.

2.2.4.9.7 Manually-Triggered Auto-Bandwidth Adjustment

Manually-triggered auto-bandwidth adjustment feature is configured with the **tools>perform>router>mpls adjust-autobandwidth [lsp lsp-name [force [bandwidth mbps]]]** command to attempt immediate auto-bandwidth adjustment for either one specific LSP or all active LSPs. If the LSP is not specified then the system assumes the command applies to all LSPs. If an LSP name is provided then the command applies to that specific LSP only and the optional **force** parameter (with or without a bandwidth) can be used.

If **force** is not specified (or the command is not LSP-specific) then measured_bw is compared to current_bw and bandwidth adjustment may or may not occur

If **force** is specified and a bandwidth is not provided then the threshold checking is bypassed but the min and max bandwidth constraints are still enforced.

If **force** is specified with a bandwidth (in Mb/s) then `signaled_bw` is set to this bandwidth. There is no requirement that the bandwidth entered as part of the command fall within the range of min-bandwidth to max-bandwidth.

The `adjust-interval`, maximum average data rate and overflow count are not reset by the manual auto-bandwidth command, whether or not the bandwidth adjustment succeeds or fails. The overflow count is reset only if the manual auto-bandwidth adjustment is successful.

2.3 RSVP

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

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

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

RSVP uses two message types to set up LSPs, PATH and RESV. [Figure 5](#) 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 FEC for which label bindings are desired. PATH messages are used to signal and request label bindings required to establish the LSP from ingress to egress. Each router along the path observes the traffic type.
PATH messages facilitate the routers along the path to make the necessary bandwidth reservations and distribute the label binding to the router upstream.
- The ELER sends label binding information in the RESV messages in response to PATH messages received.

- The LSP is considered operational when the ILER receives the label binding information.



Figure 5 Establishing LSPs**Figure 6 LSP Using RSVP Path Set Up**

Figure 6 displays an example of an LSP path set up using RSVP. The ingress label edge router (ILER 1) transmits an RSVP path message (path: 30.30.30.1) downstream to the egress label edge router (ELER 4). The path message contains a label request object that requests intermediate LSRs and the ELER to provide a label binding for this path.

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

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

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

2.3.1 Using RSVP for MPLS

Hosts and routers that support both MPLS and RSVP can associate labels with RSVP flows. When MPLS and RSVP are combined, the definition of a flow can be made more flexible. 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 FEC which defines the RSVP flow.

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

2.3.1.1 RSVP Traffic Engineering Extensions for MPLS

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

2.3.1.2 Hello Protocol

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

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

2.3.1.3 MD5 Authentication of RSVP Interface

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

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

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

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

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

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

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

2.3.1.4 Configuring Authentication using Keychains

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

Within RSVP, authentication must be explicitly configured through the use of the authentication keychain mechanism. This mechanism allows for the configuration of authentication keys and allows the keys to be changed without affecting the state of the protocol adjacencies.

To configure the use of an authentication keychain within RSVP, use the following steps:

1. Configure an authentication keychain within the `config>system>security` context. The configured keychain must include at least one valid key entry, using a valid authentication algorithm for the RSVP protocol.
2. Associate the configured authentication keychain with RSVP at the interface level of the CLI, this is done through the use of the `"auth-keychain name"` 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 in the key entry must be supported by the RSVP protocol.

The RSVP protocol supports the following algorithms:

- clear text password
- HMAC-MD5
- HMAC-SHA-1

Error handling:

- If a keychain exists but there are no active key entries with an authentication type that is valid for the associated protocol then inbound protocol packets will not be authenticated and discarded, and no outbound protocol packets should be sent.
- If keychain exists but the last key entry has expired, a log entry will be raised indicating that all keychain entries have expired. The RSVP protocol requires that the protocol not revert to an unauthenticated state and requires that the old key is not to be used, therefore, once the last key has expired, all traffic will be discarded.

2.3.2 Reservation Styles

LSPs can be signaled with explicit reservation styles. A reservation style is a set of control options that specify a number of supported parameters. The style information is part of the LSP configuration. SR OS 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 FRR option is enabled for the LSP and selects the facility FRR method at the head-end node, only the SE reservation style is allowed. Furthermore, if a PLR node receives a path message with fast-reroute requested with facility method and the FF reservation style, it will reject the reservation. The one-to-one detour method supports both FF and SE styles.

2.3.2.1 RSVP Message Pacing

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

2.3.3 RSVP Overhead Refresh Reduction

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

- RSVP message bundling — This capability is intended to reduce overall message handling load. The system supports receipt and processing of bundled message only, but no transmission of bundled messages.
- Reliable message delivery — This capability consists of sending a message-id and returning a message-ack for each RSVP message. It can be used to detect message loss and support reliable RSVP message delivery on a per hop basis. It also helps reduce the refresh rate since the delivery becomes more reliable.
- Summary refresh — This capability consists of refreshing multiples states with a single message-id list and sending negative ACKs (NACKs) for a message_id which could not be matched. The summary refresh capability reduce the amount of messaging exchanged and the corresponding message processing between peers. It does not however reduce the amount of soft state to be stored in the node.

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

The RSVP Overhead Refresh Reduction is supported with both RSVP P2P LSP path and the S2L path of an RSVP P2MP LSP instance over the same RSVP interface.

2.3.4 RSVP Graceful Restart Helper

The **gr-helper** command enables the RSVP Graceful Restart Helper feature.

The RSVP-TE Graceful Restart helper mode allows the SR OS based system (the helper node) to provide another router that has requested it (the restarting node) a grace period, during which the system will continue to use RSVP sessions to neighbors requesting the grace period. This is typically used when another router is rebooting its control plane but its forwarding plane is expected to continue to forward traffic based on the previously available Path and Resv states.

The user can enable Graceful Restart helper on each RSVP interface separately. When the GR helper feature is enabled on an RSVP interface, the node starts inserting a new Restart_Cap Object in the Hello packets to its neighbor. The restarting node does the same and indicates to the helper node the desired Restart Time and Recovery Time.

The GR Restart helper consists of a couple of phases. Once it loses Hello communication with its neighbor, the helper node enters the Restart phase. During this phase, it preserves the state of all RSVP sessions to its neighbor and waits for a new Hello message.

Once the Hello message is received indicating the restarting node preserved state, the helper node enters the recovery phase in which it starts refreshing all the sessions that were preserved. The restarting node will activate all the stale sessions that are refreshed by the helper node. Any Path state that did not get a Resv message from the restarting node once the Recovery Phase time is over is considered to have expired and is deleted by the helper node causing the proper Path Tear generation downstream.

The duration of the restart phase (recovery phase) is equal to the minimum of the neighbor's advertised Restart Time (Recovery Time) in its last Hello message and the locally configured value of the max-restart (max-recovery) parameter.

When GR helper is enabled on an RSVP interface, its procedures apply to the state of both P2P and P2MP RSVP LSP to a neighbor over this interface.

2.3.5 Enhancements to RSVP Control Plane Congestion Control

The RSVP control plane makes use of a global flow control mechanism to adjust the rate of Path messages for unmapped LSP paths sent to the network under congestion conditions. When a Path message for establishing a new LSP path or retrying an LSP path that failed is sent out, the control plane keeps track of the rate of successful establishment of these paths and adjusts the number of Path messages it sends per second to reflect the success ratio.

In addition, an option to enable an exponential back-off retry-timer is available. When an LSP path establishment attempt fails, the path is put into retry procedures and a new attempt will be performed at the expiry of the user-configurable retry-timer. By default, the retry time is constant. The exponential back-off timer procedures will double the value of the user configurable retry-timer value at every failure of the attempt to adjust to the potential network congestion that caused the failure. An LSP establishment fails if no Resv message was received and the Path message retry-timer expired, or a PathErr message was received before the timer expired.

Three enhancements to this flow-control mechanism to improve congestion handling in the rest of the network are supported.

The first enhancement is the change to the LSP path retry procedure. If the establishment attempt failed due to a Path message timeout and no Resv was received, the next attempt will be performed at the expiry of a new LSP path initial retry-timer instead of the existing retry-timer. While the LSP path initial retry-timer is still running, a refresh of the Path message using the same path and the same LSP-id is performed according to the configuration of the refresh-timer. Once the LSP path initial retry-timer expires, the ingress LER then puts this path on the regular retry-timer to schedule the next path signaling using a new computed path by CSPF and a new LSP-id.

The benefits of this enhancement is that the user can now control how many refreshes of the pending PATH state can be performed before starting a new retry-cycle with a new LSP-id. This is all done without affecting the ability to react faster to failures of the LSP path, which will continue to be governed by the existing retry-timer. By configuring the LSP path initial retry-timer to values that are larger than the retry-timer, the ingress LER will decrease the probability of overwhelming a congested LSR with new state while the previous states installed by the same LSP are lingering and will only be removed after the refresh timeout period expires.

The second enhancement consists of applying a jitter +/- 25% to the value of the retry-timer similar to how it is currently done for the refresh timer. This will further decrease the probability that ingress LER nodes synchronize their sending of Path messages during the retry-procedure in response to a congestion event in the network.

The third enhances the RSVP flow control mechanism by taking into account new parameters: outstanding CSPF requests, Resv timeouts and Path timeouts.

2.3.6 RSVP LSP Statistics

This feature provides the following counters:

- 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

The counters are available for an RSVP LSP at the egress datapath of an ingress LER and at the ingress datapath of an egress LER. No LSR statistics are provided.

This feature is supported on IOM-3 and requires chassis mode D.

2.3.7 P2MP RSVP-TE LSP Statistics

This feature provides the following counters for a RSVP P2MP LSP instance:

- 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.

The above counters are provided for the following LSR roles:

- At ingress LER, a set of per P2MP LSP instance counters for packets forwarded to the P2MP LSP instance without counting the replications is provided. In other words, a packet replicated over multiple branches of the same P2MP LSP instance will count once as long as at least one LSP branch forwarded it.

- At BUD LSR and egress LER, per ILM statistics are provided. These counters will include all packets received on the ILM, whether they match a L2/L3 MFIB record or not. ILM stats will work the same way as for a P2P LSP. In other words, they will count all packets received on the primary ILM, including packets received over the bypass LSP.

When MBB is occurring for an S2L path of an RSVP P2MP LSP, paths of the new and old S2L will both receive packets on the egress LER. Both packets are forwarded to the fabric and outgoing PIM/IGMP interfaces until the older path is torn down by the ingress LER. In this case, packet duplication should be counted.

- No branch LSR statistics are provided.
- The P2MP LSP statistics share the same pool of counters and stat indices the P2P LSP share on the node. Each P2P/P2MP RSVP LSP or LDP FEC consumes one stat index for egress stats and one stat index for ingress stats.
- The user can retrieve the above counters in four different ways:
 - In CLI display of the output of the show command applied to a specific instance, or a specific template instance, of a RSVP P2MP.
 - In CLI display of the output of the monitor command applied to a specific instance, or a specific template instance, of a RSVP P2MP.
 - Via an SNMP interface by querying the MIB.
 - Via an accounting file if statistics collection with the default or user specified accounting policy is enabled for the MPLS LSP stats configuration contexts.
- OAM packets that are forwarded using the LSP encapsulation, for example, P2MP LSP Ping and P2MP LSP Trace, are also included in the above counters.

The user can determine if packets are dropped for a given branch of a P2MP RSVP LSP by comparing the egress counters at the ingress LER with the ILM counters at the egress LER or BUD LSR.

Octet counters are for the entire frame and thus include the label stack and the L2 header and padding similar to the existing P2P RSVP LSP and LDP FEC counters. Thus ingress and egress octet counters for an LSP may slightly differ if the type of interface or encapsulation is different (POS, Ethernet NULL, Ethernet Dot1.Q).

2.3.7.1 Configuring RSVP P2MP LSP Egress Statistics

At ingress LER, the configuration of the egress statistics is under the MPLS P2MP LSP context when carrying multicast packets over a RSVP P2MP LSP in the base routing instance. This is the same configuration as the one already supported with P2P RSVP LSP.

```

config
router
  [no] mpls
    [no] lsp lsp-name p2mp-lsp
      [no] egress-statistics
        accounting-policy policy-id
        no accounting-policy
      [no] collect-stats
    [no] shutdown

```

If there are no stat indices available when the user performs the 'no shutdown' command for the egress statistics node, the command will be failed.

The configuration is in the P2MP LSP template when the RSVP P2MP LSP is used as an I-PMSI or S-PMSI in multicast VPN or in VPLS/B-VPLS.

```

config
router
  [no] mpls
    lsp-template template-name p2mp
    no lsp-template template-name
      [no] egress-statistics
        accounting-policy policy-id
        no accounting-policy
      [no] collect-stats

```

If there are no stat indices available at the time an instance of the P2MP LSP template is signaled, no stats are allocated to the instance, but the LSP is brought up. In this case, an operational state of out-of-resources is shown for the egress stats in the show output of the P2MP LSP S2L path.

2.3.7.2 Configuring RSVP P2MP LSP Ingress Statistics

When the ingress LER signals the path of the S2L sub-LSP, it includes the name of the LSP and that of the path in the Session Name field of the Session Attribute object in the Path message. The encoding is as follows:

Session Name: *lsp-name::path-name*, where *lsp-name* component is encoded as follows:

1. P2MP LSP via user configuration for L3 multicast in global routing instance: "LspNameFromConfig"
2. P2MP LSP as I-PMSI or S-PMSI in L3 mVPN: *templateName-SvcId-mTTmIndex*
3. P2MP LSP as I-PMSI in VPLS/B-VPLS: *templateName-SvcId-mTTmIndex*

The ingress statistics CLI configuration allows the user to match either on the exact name of the P2MP LSP as configured at the ingress LER or on a context which matches on the template name and the service-id as configured at the ingress LER.

```
config
  router
    [no] mpls
      ingress-statistics
        [no] lsp lsp-name sender sender-address
            accounting-policy policy-id
            no accounting-policy
        [no] collect-stats
        [no] shutdown

        [no] p2mp-template-lsp rsvp-session-name
            SessionNameString sender sender-address
            accounting-policy policy-id
            no accounting-policy
        [no] collect-stats
            max-stats integer<1-8192 | max, default max>
            no max-stats
        [no] shutdown
```

When the matching is performed on a context, the user must enter the RSVP session name string in the format "*templateName-svcId*" to include the LSP template name as well as the mVPN VPLS/B-VPLS service ID as configured at the ingress LER. In this case, one or more P2MP LSP instances signaled by the same ingress LER could be associated with the ingress statistics configuration. In this case, the user is provided with CLI parameter **max-stats** to limit the maximum number of stat indices which can be assigned to this context. If the context matches more than this value, the additional request for stat indices from this context will be rejected.

The rules when configuring an ingress statistics context based on template matching are the following:

1. **max-stats** once allocated can be increased but not decreased unless the entire ingress statistics context matching a template name is deleted.
2. In order to delete ingress statistics context matching a template name, a shutdown is required.
3. An accounting policy cannot be configured or de-configured until the ingress statistics context matching a template name is shutdown.
4. After deleting an accounting policy from an ingress statistics context matching a template name, the policy is not removed from the log until a 'no shut' is performed on the ingress statistics context.

If there are no stat indices available at the time the session of the P2MP LSP matching a template context is signaled and the session state installed by the egress LER, no stats are allocated to the session.

Furthermore, the assignment of stat indices to the LSP names that match the context will also be not deterministic. The latter is due to the fact that a stat index is assigned and released following the dynamics of the LSP creation or deletion by the ingress LER. For example, a multicast stream crosses the rate threshold and is moved to a newly signaled S-PMSI dedicated to this stream. Later on, the same stream crosses the threshold downwards and is moved back to the shared I-PMSI and the P2MP LSP corresponding to the S-PMSI is deleted by the ingress LER.

2.3.7.3 Configuring Implicit Null

The implicit null label option allows a router egress LER to receive MPLS packets from the previous hop without the outer LSP label. The operation of the previous hop is referred to as penultimate hop popping (PHP).

This option is signaled by the egress LER to the previous hop during the LSP signaling with RSVP control protocol. In addition, the egress LER can be configured to receive MPLS packet with the implicit null label on a static LSP.

The user can configure your router to signal the implicit null label value over all RSVP interfaces and for all RSVP LSPs for which this node is the egress LER using the **implicit-null-label** command in the **config>router>rsvp** context.

The user must shutdown RSVP before being able to change the implicit null configuration option.

The user can also override the RSVP level configuration for a specific RSVP interface:

```
config>router>rsvp>if>implicit-null-label {enable | disable}
```

All LSPs for which this node is the egress LER and for which the path message is received from the previous hop node over this RSVP interface will signal the implicit null label. This means that if the egress LER is also the merge-point (MP) node, then the incoming interface for the path refresh message over the bypass dictates if the packet will use the implicit null label or not. The same for a 1-to-1 detour LSP.

By default, an RSVP interface inherits the RSVP level configuration. The user must shutdown the RSVP interface before being able to change the implicit null configuration option.



Note: The RSVP interface must be shutdown regardless of whether the new value for the interface is the same or different than the one it is currently using.

The egress LER does not signal the implicit null label value on P2MP RSVP LSPs. However, the PHP node can honor a Resv message with the label value set to the implicit null value when the egress LER is a third party implementation.

The implicit null label option is also supported on a static label LSP. The following commands can be used to cause the node to push or to swap to an implicit null label on the MPLS packet:

```
config>router>mpls>static-lsp>push implicit-null-label nexthop ip-address
```

```
config>router>mpls>if>label-map>swap implicit-null-label nexthop ip-address
```

2.3.8 Using Unnumbered Point-to-Point Interface in RSVP

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

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

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

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

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

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

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

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

The unnumbered IP is advertised by IS-IS TE and OSPF TE, and CSPF can include them in the computation of a path for a P2P LSP or for the S2L of a P2MP LSP. This feature does not, however, support defining an unnumbered interface 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 a given unnumbered interface index is obtained from the TE database. As a result, if traffic engineering is disabled in IS-IS or OSPF, a non-CSPF LSP with the next-hop for its path is over an unnumbered interface will not come up at the ingress LER since the router-id of the neighbor that has the next-hop of the path message cannot be looked up. In this case, the LSP path will remain in operationally down state with a reason noRouteToDestination. If a PATH message was received at the LSR in which traffic engineering was disabled and the next-hop for the LSP path is over an unnumbered interface, a PathErr message will be sent back to the ingress LER with the *Routing Problem* error code of 24 and an error value of 5 "No route available toward destination".

All MPLS features available for numbered IP interfaces are supported, with the exception of the following:

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

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

2.3.8.1 Operation of RSVP FRR Facility Backup over Unnumbered Interface

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

The router uses the address of the outgoing interface of the bypass LSP as the *IPv4 tunnel sender address* in the sender template object. This address will be different from the system interface address used in the sender template of the protected LSP by the ingress LER and thus there are no conflicts when the ingress LER acts as a PLR.

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

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

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

- Change in C-Type of the RSVP_HOP object, or
- C-Type is IF_ID RSVP_HOP and did not change but IF_ID TLV is different, or
- Change in IPv4 Next/Previous Hop Address in RSVP_HOP object regardless of the C-Type value.

These procedures at PLR and MP nodes are followed in both link-protect and node-protect FRR. If the MP node is running a pre-Release 11.0 implementation, it will reject the new IF_ID C-Type and will drop the PATH over bypass. This will result in the protected LSP state expiring at the MP node, which will tear down the path. This will be the case in general when node-protect FRR is enabled and the MP node does not support unnumbered RSVP interface.

2.4 MPLS Transport Profile

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

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

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

In order to meet these objectives, MPLS-TP has a number of high level characteristics:

- It does not modify the MPLS forwarding architecture, which is based on existing pseudowire and LSP constructs. Point-to-point LSPs may be unidirectional or bi-directional. Bi-directional LSPs must be congruent (that is, co-routed and follow the same path in each direction). The system supports bidirectional co-routed MPLS-TP LSPs.
- There is no LSP merging.
- OAM, protection, and forwarding of data packets can operate without IP forwarding support. When static provisioning is used, there is no dependency on dynamic routing or signaling.
- LSP and pseudowire monitoring is only achieved through the use of OAM and does not rely on control plane or routing functions to determine the health of a path. For example, LDP hello failures do not trigger protection.
- MPLS-TP can operate in the absence of an IP control plane and IP forwarding of OAM traffic. In release 11.0, MPLS-TP is only supported on static LSPs and PWs.

The system supports MPLS-TP on LSPs and PWs with static labels. MPLS-TP is not supported on dynamically signaled LSPs and PWs. MPLS-TP is supported for Epipe, Apipe, and Cpipe VLLs, and Epipe spoke-SDP termination on IES, VPRN and VPLS. Static PWs may use SDPs that use either static MPLS-TP LSPs or RSVP-TE LSPs.

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


- MPLS-TP Generic Associated Channel for LSPs and PWs (RFC 5586)
- MPLS-TP Identifiers (RFC 6370)
- Proactive CC, CV, and RDI using BFD for LSPs (RFC 6428)
- On-Demand CV for LSPs and PWs using LSP Ping and LSP Trace (RFC 6426)
- 1-for-1 Linear protection for LSPs (RFC 6378)
- Static PW Status Signaling (RFC 6478)

The system can play the role of an LER and an LSR for static MPLS-TP LSPs, and a PE/T-PE and an S-PE for static MPLS-TP PWs. It can also act as a S-PE for MPLS-TP segments between an MPLS network that strictly follows the transport profile, and an MPLS network that supports both MPLS-TP and dynamic IP/MPLS.

2.4.1 MPLS-TP Model

Figure 7 shows a high level functional model for MPLS-TP in SR OS. LSP A and LSP B are the working and protect LSPs of an LSP tunnel. These are modeled as working and protect paths of an MPLS-TP LSP in SR OS. MPLS-TP OAM runs in-band on each path. 1:1 linear protection coordinates the working and protect paths, using a protection switching coordination protocol (PSC) that runs in-band on each path over a Generic Associated Channel (G-ACh) on each path. Each path can use either an IP numbered, IP unnumbered, or MPLS-TP unnumbered (that is, non-IP) interface.

Figure 7 MPLS-TP Model



al_0221

All MPLS-TP LSPs are bidirectional co-routed, as detailed in RFC5654. That is, the forward and backward directions follow the same route (in terms of links and nodes) across the network. Both directions are set up, monitored and protected as a single entity. Therefore, both ingress and egress directions of the same LSP segment are associated at the LER and LSR and use the same interface (although this is not enforced by the system).

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

2.4.2 MPLS-TP Provider Edge and Gateway


This section describes some example roles for the system in an MPLS-TP network.

2.4.2.1 VLL Services

The system may use MPLS TP LSPs, and PWs, to transport point to point virtual leased line services. The router may play the role of a terminating PE or switching PE for VLLs. Epipe, Apipe, and Cpipe VLLs are supported.



[Figure 8](#) illustrates the use of the system as a T-PE for services in an MPLS-TP domain, and as a S-PE for services between an MPLS-TP domain and an IP/MPLS domain. Static PWs with MPLS-TP identifiers, originating in the MPLS-TP network, are transported over static MPLS-TP LSPs. These either terminate on a local SAP on the system, or are switched to another PW segment across the IP/MPLS network. The PW segment in the IP/MPLS network may have static labels or be signaled using T-LDP.

Figure 8 MPLS-TP Provider Edge and Gateway, VLL Services



2.4.2.2 Spoke-SDP Termination


Figure 9 and Figure 10 illustrate the model for spoke-sdp termination on VPLS and IES/VP RN services, respectively. Similar to the VLL case, the static MPLS-TP PW may terminate on an interface belonging to the service on the router at the border between the MPLS-TP and IP/MPLS networks, or be switched to another PW segment to be terminated on a remote PE.

Figure 9 MPLS-TP Provider Edge and Gateway, spoke-SDP Termination on VPLS**Figure 10 MPLS-TP Provider Edge and Gateway, spoke-SDP Termination on IES/VPRN**

2.4.3 MPLS-TP LSR

The SR OS MPLS-TP LSR model is illustrated in [MPLS-TP LSR](#). The system is able to swap a statically configured LSP label on an ingress path to a statically configured LSP label on an egress path. Bidirectional co-routed MPLS TP LSPs are supported by configuring the forward and reverse paths of the LSP to use the same ports on ingress and egress.

Figure 11 MPLS-TP LSR



2.4.4 Detailed Descriptions of MPLS-TP

2.4.4.1 MPLS-TP LSPs

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

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

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

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

The router may play the role of either an LER or an LSR.

2.4.4.2 MPLS-TP on Pseudowires

MPLS-TP is supported on PWs with static labels. The provisioning model supports RFC6370-style PW path identifiers for MPLS-TP PWs.

MPLS-TP PWs reuse the static PW provisioning model of previous SR OS releases. Including the use of the PW-switching key work to distinguish an S-PE. Therefore, the primary distinguishing feature for an MPLS-TP PW is the ability to configure MPLS-TP PW path identifiers, and to support MPLS-TP OAM and static PW status signaling.

The system can perform the role of a T-PE or an S-PE for a PW with MPLS-TP.


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

2.4.5 MPLS-TP Maintenance Identifiers

MPLS-TP is designed for use both with, and without, a control plane. MPLS-TP therefore specifies a set of identifiers that can be used for objects in either environment. This includes a path and maintenance identifier architecture comprising Node, Interface, PW and LSP identifiers, Maintenance Entity Groups (MEGs), Maintenance End Points (MEPs) and Maintenance Intermediate Points (MIPs). These identifiers are specified in RFC6370.

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

Figure 12 MPLS-TP Maintenance Architecture



al_0226

Both IP-compatible and ICC (ITU-T carrier code) based identifiers for the above objects are specified in the IETF, but only the IP-compatible identifiers defined in RFC6370 are supported.

SR OS supports the configuration of the following node and interface related identifiers:

- **Global_ID:** this is similar to the global ID that can be configured for Dynamic MS-PWs in Release 9.0 of SR OS. However, in MPLS-TP this should be set to the AS# of the node. If not explicitly configured, then it assumes the default value of 0. In SR OS, the source Global ID for an MPLS-TP Tunnel is taken to be the Global ID configured at the LER. The destination Global ID is optional in the tunnel configuration. If it is not configured, then it is taken as the same as the source Global ID.
- **Node_ID:** This is a 32-bit value assigned by the operator within the scope of the Global_ID. The system supports the configuration of an IPv4 formatted address <a.b.c.d> or an unsigned 32-bit integer for the MPLS-TP Node ID at each node. The node ID must be unique within the scope of the global ID, but there is no requirement for it to be a valid routable IP address. Indeed, a node-id can represent a separate IP-compatible addressing space that may be separate

from the IP addressing plan of the underlying network. If no node ID is configured, then the node ID is taken to be the system interface IPv4 address of the node. When configuring a tunnel at an LER, either an IPv4 or an unsigned integer Node ID can be configured as the source and destination identifiers, but both ends must be of the same type.


- **IF_ID:** This is an MPLS-TP section layer identifier at the MPLS interface level. On the router, this is used to provide an identifier for the LSP-Trace DMAP when an IP identifier is not available. The IF_ID is a 64-bit identifier of an MPLS-TP interface on a node that is unique within the scope of a Global_ID. It is composed of the Node_ID and the IF_Num. The IF_Num is a node-wide unique identifier for an MPLS-TP interface. On the router, this is primarily used for supporting the DMAP TLV in LSP Trace using MPLS-TP identifiers with unnumbered MPLS-TP interfaces.

Statically configured LSPs are identified using GMPLS-compatible identifiers with the addition of a Tunnel_Num and LSP_Num. As in RSVP-TE, tunnels represent, for example, a set of working and protect LSPs. These are GMPLS-compatible because GMPLS chosen by the IETF as the control plane for MPLS-TP LSPs, although this is not supported in Release 11.0 of the software. PWs are identified using a PW Path ID which has the same structure as FEC129 All Type 2.

SR OS derives the identifiers for MEPs and MIPs on LSPs and PWs based on the configured identifiers for the MPLS-TP Tunnel, LSP or PW Path ID, for use in MPLS-TP OAM and protection switching, as per RFC6370.

The information models for LSPs and PWs supported in Release 11.0 are illustrated in [Figure 13](#) and [Figure 14](#). The figures use the terminology defined in RFC6370.


Figure 13 MPLS-TP LSP and Tunnel Information Model



The MPLS-TP Tunnel ID and LSP ID are not to be confused with the RSVP-TE tunnel id implemented on the router system. Table 4 shows how these map to the X and Y ends of the tunnel shown in Figure 13 for the case of co-routed bidirectional LSPs.

Table 4 Mapping from RSVP-TE to MPLS-TP Maintenance Identifiers

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

Figure 14 MPLS-TP PW Information Model


al_0228

In the PW information model shown in Figure 14, the MS-PW is identified by the PW Path ID that is composed of the full AGI:SAIL:TAIL. The PW Path ID is also the MEP ID at the T-PEs, so a user does not have to explicitly configure a MEP ID; it is automatically derived by the system. For MPLS-TP PWs with static labels, although the PW is not signaled end-to-end, the directionality of the SAIL and TAIL is taken to be the same as for the equivalent label mapping message i.e. from downstream to upstream. This is to maintain consistency with signaled pseudowires using FEC 129.

On the system, an S-PE for an MS-PW with static labels is configured as a pair of spoke-sdps bound together in an VLL service using the VC-switching command. Therefore, the PW Path ID configured at the spoke-SDP level at an S-PE must contain the Global-ID, Node-ID and AC-ID at the far end T-PEs, not the local S-PE. The ordering of the SAIL:TAIL in the PW Path ID where static PWs are used should be consistent with the direction of signaling of the egress label to a spoke-SDP forming that segment, if that label were signaled using T-LDP (in downstream unsolicited mode). VCCV Ping will check the PW ID in the VCCV Ping echo request message against the configured PW Path ID for the egress PW segment.

Figure 15 shows an example of how the PW Path IDs can be configured for a simple two-segment MS-PW.

Figure 15 Example usage of PW Identifiers




2.4.5.1 Generic Associated Channel

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

RFC5586 specifies mechanisms for implementing the G-ACh, relying on the combination of a reserved MPLS label, the 'Generic-ACH Label (GAL)', as an alert mechanism (value=13) and Generic Associated Channel Header (G-ACh) for MPLS LSPs, and using the Generic Associated Channel Header, only, for MPLS PWs (although the GAL is allowed on PWs). The purpose of the GAL is to indicate that a G-ACh resides at the bottom of the label stack, and is only visible when the bottom non-reserved label is popped. The G-ACh channel type is used to indicate the packet type carried on the G-ACh. Packets on a G-ACh are targeted to a node containing a MEP by ensuring that the GAL is pushed immediately below the label

that is popped at the MEP (e.g. LSP endpoint or PW endpoint), so that it can be inspected as soon as the label is popped. A G-ACh packet is targeted to a node containing a MIP by setting the TTL of the LSP or PW label, as applicable, so that it expires at that node, in a similar manner to the SR OS implementation of VCCV for MS-PWs.

Figure 16 Label for LSP and PW G-ACh Packets



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

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

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

2.4.5.2.1 On-Demand Connectivity Verification (CV) using LSP-Ping

MPLS-TP supports mechanisms for on demand CC/CV as well as route tracing for LSPs and PWs. These are required to enable an operator to test the initial configuration of a transport path, or to assist with fault isolation and diagnosis. On demand CC/CV and route tracing for MPLS-TP is based on LSP-Ping and is described in RFC6426. Three possible encapsulations are specified in that RFC:

- IP encapsulation, using the same label stack as RFC4379, or encapsulated in the IPv4 G-ACh channel with a GAL/ACH
- and non-IP encapsulation with GAL/ACH for LSPs and ACH for PWs.

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

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

The routers support the following encapsulations:

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

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

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

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

Downstream Mapping Support

RFC 4379 specifies four address types for the downstream mapping TLV for use with IP numbered and unnumbered interfaces, as listed in [Table 5](#):

Table 5 Downstream Mapping (RFC 4379)

Type #	Address Type	K Octets	Reference
1	IPv4 Numbered	16	RFC 4379
2	IPv4 Unnumbered	16	
3	IPv6 Numbered	40	
4	IPv6 Unnumbered	28	

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

It is possible to send and respond to a DSMAP/DDMAP TLV in the LSP Trace packet for numbered IP interfaces as per RFC4379. In this case, the echo request message contains a downstream mapping TLV with address type 1 (IPv4 address) and the IPv4 address in the DDMAP/DSMAP TLV is taken to be the IP address of the IP interface that the LSP uses. The LSP trace packet therefore contains a DSMAP TLV in addition to the MPLS-TP static LSP TLV in the target FEC stack.

DSMAP/DDMAP is not supported for pseudo wires.

2.4.5.2.2 Proactive CC, CV and RDI

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

Proactive CC/CV based on bidirectional forwarding detection (BFD) for MPLS-TP is described in RFC6428. BFD packets are sent using operator configurable timers and encapsulated without UDP/IP headers on a standardized G-ACh channel on an LSP or PW. CC packets simply consist of a BFD control packet, while CV packets also include an identifier for the source MEP in order that the sink MEP can detect if it is receiving packets from an incorrect peer MEP, thus indicating a mis-connectivity defect. Other defect types (including period mis-configuration defect) should be

supported. When a supported defect is detected, an appropriate alarm is generated (e.g. log, SNMP trap) at the receiving MEP and all traffic on the associated transport path (LSP or PW) is blocked. This is achieved using linear protection for CC defects, and by blocking the ingress data path for CV defects. The system supports both a CC-only mode and a combine CC / CV mode, as defined in RFC6428.

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

Figure 17 BFD used for proactive CC on MPLS-TP LSP





Figure 18 BFD used for proactive CV on MPLS-TP LSP



Linear protection switching of LSPs (see below) is triggered based on a CC or CV defect detected by BFD CC/CV.

RFC6428 defines two BFD session modes: Coordinated mode, in which the session state on both directions of the LSP is coordinated and constructed from a single, bidirectional BFD session, and independent mode, in which two independent sessions are bound together at a MEP. Coordinated mode is supported.

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

The following GACH channel types are supported for the combined CC/CV mode:

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

The following G-ACh channel types are used for the CC-only mode:

- 0x07

2.4.5.2.3 BFD-based RDI

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

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

2.4.5.3 PW Control Channel Status Notifications (Static Pseudowire Status Signaling)

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

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

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

The system supports the optional acknowledgment of a PW control channel status message.

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

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

A trap is generated if the refresh timer times-out.

If PW redundancy is configured, the system will always consider the literal value of the PW status; a time-out of the refresh timer will not impact the choice of the active transit object for the VLL service. The result of this is that if the refresh timer times-out, and a given PW is currently the active PW, then the system will not fail-over to an alternative PW if the status is zero and some lower-layer OAM mechanism e.g. BFD has not brought down the LSP due to a connectivity defect. It is recommended that the PW refresh timer be configured with a much longer interval than any proactive OAM on the LSP tunnel, so that the tunnel can be brought down before the refresh timer expires if there is a CC defect.

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

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

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

2.4.5.4 PW Control Channel Status Request Mechanism

The system implements an optional PW control channel status request mechanism. This enhances the existing control channel status mechanism so that a peer that has *stale* PW status for the far-end of a PW can request that the peer PE send a static PW status update. Accurate and current information about the far end status of a PW is important for proper operation of PW redundancy. This mechanism ensures a consistent view of the control plane is maintained, as far as possible, between peer nodes. It is not intended to act as a continuity check between peer nodes.

2.4.5.5 Pseudowire Redundancy and Active / Standby Dual Homing

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

The following PW redundancy scenarios must be supported:

- MC-LAG and MC-APS with single and multi-segment PWs interconnecting the PEs.
- MS-PW (S-PE) Redundancy between VLL PEs with single-homed CEs.

- Dual-homing of a VLL service into redundant IES or VPRN PEs, with active/standby PWs.
- Dual-homing of a VLL service into a VPLS with active/standby PWs.

Active/standby dual-homing into routed VPLS is not supported in for MPLS-TP PWs. This is because it relies on PW label withdrawal of the standby PW in order to take down the VPLS instance, and hence the associated IP interface. Instead, it is possible to enable BGP multi-homing on a routed VPLS that has MPLS-TP PWs as spokes, and for the PW status of each spoke-SDP to be driven (using control channel status) from the active or standby forwarding state assigned to each PW by BGP.

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

A failover to a standby pseudowire is initiated based on the existing supported methods (e.g. failure of the SDP).

2.4.5.6 Lock Instruct and Loopback for MPLS-TP Pseudowires

On the 7750 SR and 7450 ESS, the MPLS-TP supports lock instruct and loopback for PWs, including the ability to:

- administratively lock a spoke-sdp with MPLS-TP identifiers
- divert traffic to and from an external device connected to a SAP
- create a data path loopback on the corresponding PW at a downstream S-PE or T-PE that was not originally bound to the spoke-sdp being tested
- forward test traffic from an external test generator into an administratively locked PW, while simultaneously blocking the forwarding of user service traffic

MPLS-TP provides the ability to conduct test service throughput for PWs, through the configuration of a loopback on an administratively locked pseudowire. To conduct a service throughput test, an administrative lock is applied at each end of the PW. A test service that contains the SAP connected to the external device is used to inject test traffic into the PW. Lock request messaging is not supported.

A lock can be applied using the CLI or NMS. The forwarding state of the PW can be either active or standby.

After the PW is locked it can be put into loopback mode (for two way tests) so the ingress data path in the forward direction is cross connected to the egress data path in the reverse direction of the PW. The loopback can be configured through the CLI or NMS.

The PW loopback is created at the PW level, so everything under the PW label is looped back. This distinguishes a PW loopback from a service loopback, where only the native service packets are looped back.

The following MPLS-TP loopback configuration is supported:

- An MPLS-TP loopback can be created for an epipe, cpipe or a pipe VLL.
- Test traffic can be inserted at an epipe, cpipe or a pipe VLL endpoint or at an epipe spoke-sdp termination on a VPLS interface.

For more information about configuring lock instruct and loopback for MPLS-TP Pseudowires see, the 7450 ESS, 7750 SR, and 7950 XRS Services Overview Guide and the 7450 ESS, 7750 SR, and 7950 XRS Services Overview Guide.

2.4.5.7 MPLS-TP LSP Protection

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

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

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


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

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

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


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

Figure 19 Normal Operation



al_0233

Figure 20 Failed Condition



al_0234

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

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

Figure 21 Failed Condition - Switching at A





Figure 22 Failed Condition - Switching at Z



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

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

If BFD CC or CC/CV OAM packets are used to detect defects on the working and protection paths, they are inserted on both working and protection paths. Packets are sent whether or not the path is selected as the currently active path. Linear protection switching is also triggered on receipt of an AIS with the LDI bit set.


The following operator commands are supported:

- Forced Switch
- Manual Switch
- Clear

2.4.6 Switching Static MPLS-TP to Dynamic T-LDP Signaled PWs

Some use cases for MPLS-TP require an MPLS-TP based aggregation network and an IP-based core network to interoperate, so providing the seamless transport of packet services across static MPLS-TP and dynamically signaled domains using an MS-PW. In this environment, end to end VCCV Ping and VCCV Trace may be used on the MS-PW. This is illustrated in [Figure 23](#):

Figure 23 Static - Dynamic PW Switching with MPLS-TP



Services are backhauled from the static MPLS-TP network on the left to the dynamic IP/MPLS network on the right. The router acts as an S-PE interconnecting the static and dynamic domains.

The router implementation supports such use cases through the ability to mate a static MPLS-TP spoke-sdp, with a defined pw-path-id, to a FEC128 spoke-sdp. The dynamically signaled spoke-sdp must be MPLS; GRE PWs are not supported, but the T-LDP signaled PW can use any supported MPLS tunnel type (e.g. LDP, RSVP-TE, static, BGP). The control-word must be enabled on both mate spoke-sdps.

Mapping of control channel status signaling to and from T-LDP status signaling at the router S-PE is also supported.


The use of VCCV Ping and VCCV Trace on an MS-PW composed of a mix of static MPLS-TP and dynamic FEC128 segments is described in more detail in the *7450 ESS, 7750 SR, and 7950 XRS OAM and Diagnostics Guide*.

2.4.7 Alarm Indication Signal (AIS)

When a MEP at a server layer (such as a link layer with respect to a given LSP) detects a failure, the server MEP notifies a co-located client layer of the condition. The client layer then generates Alarm Indication Signal (AIS) packets downstream in the client layer. These fault OAM messages are generated by intermediate nodes where a client LSP is switched, as per RFC 6427. This means that AIS packets are only inserted at an LSP MIP. AIS is used by the receiving MEP to suppress client layer traps caused by the upstream server layer failure; for example, if BFD CC is running on the LSP, then AIS will suppress the generation of multiple traps due to loss of CC.

[Figure 24](#) illustrates an example of the operation of AIS in MPLS-TP.

Figure 24 Example of AIS in MPLS-TP



25535

In the example, a failure of the Ethernet link layer between PE1 and LSR1 is detected at LSR1, which raises a local trap. LSPs transiting the LSR may be running CC OAM, such as BFD, and have AIS packets injected into them at LSR1. These AIS messages are received by the corresponding downstream MEP and processed. The failure of the Ethernet link between PE1 and LSR1 means that CC OAM on the LSPs is not received by the MEPs at PE2. Normally, this would cause multiple traps to be raised at PE2, but the reception of AIS causes PE2 to suppress the local generation of traps related to the failed LSP.

In order for traps to be suppressed successfully, the AIS message must arrive and be processed at the far-end PE or LER in sufficient time for the initial alarm to be suppressed. The router therefore implements a 2.5 secs hold-down timer for such traps on MPLS-TP LSPs.

Fault management for MPLS-TP, including AIS, is specified in RFC 6427.

The router supports:

- receiving and processing of AIS messages at LSP MEPs (at the LER)
- generation of AIS messages at LSP MIPs (at the LSR) in response to a failure of the ingress link
- suppression of SNMP traps indicating changes in the state of a BFD session, which result from the failure of the LSP data path upstream of a receiving LER; these traps would otherwise be sent to the 5620 SAM

- suppression of any BFD state machine Up/Down changes that occur while AIS is being received; there is no buffering or storage of state machine changes that occur during this period. This suppression only applies to Up/Down state change traps; other traps that would be expected are observed as normal.
- inclusion of the Link Down Indication (LDI) in an AIS message. This triggers a switchover of LSP linear protection if used on the LSP.
- insertion of AIS in the downstream direction of the transit path if a unidirectional fault is detected at an LSR. This suppresses CC traps at the downstream LER. However, the BFD session will still go down, causing RDI to be sent upstream in BFD, which will cause an alarm at the upstream LER.

2.4.8 Configuring MPLS-TP

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

2.4.8.1 Configuration Overview

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

At the router LER and LSR:

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

At the router LER, configure:

1. OAM Templates. These contain generic parameters for MPLS-TP proactive OAM. An OAM template is configured under **config>router>mpls>mpls-tp>oam-template**.
2. BFD templates. These contain generic parameters for BFD used for MPLS-TP LSPs. A BFD template is configured under **config>router>bfd>bfd-template**.

3. Protection templates. These contain generic parameters for MPLS-TP 1-for-1 linear protection. A protection template is configured under **config>router>mpls>mpls-tp>protection-template**.
4. MPLS-TP LSPs are configured under **config>router>mpls>lsp mpls-tp**
5. Pseudowires using MPLS-TP are configured as spoke-sdps with static PW labels.

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

The following sections describe these configuration steps in more detail.

2.4.8.2 Node-Wide MPLS-TP Parameter Configuration

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

The mpls-tp context is configured as follows:

```
config
router
mpls
    [no] mpls-tp
    . . .
    [no] shutdown
```

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

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

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

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

2.4.8.3 Node-Wide MPLS-TP Identifier Configuration

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

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

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

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

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

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

2.4.8.4 Static LSP and Pseudowire (VC) Label and Tunnel Ranges

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

```
config
  router
    mpls-labels
      [no] static-label max-lsp-labels <number>
    static-svc-label <number>
```

<number>: indicates the maximum number of labels for the label type.

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

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

```
config
  router
    mpls
      mpls-tp
        [no] tp-tunnel-id-range <start-id> <end-id>
```

The tunnel ID range referred to here is a contiguous range of RSVP-TE Tunnel IDs is reserved for use by MPLS TP, and these IDs map to the MPLS-TP Tunnel Numbers. There are some cases where the dynamic LSPs may have caused fragmentation to the number space such that contiguous range {max-min} is not available. In these cases, the command will fail.

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

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

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

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

config>router>mpls>static-lsp

config>router>mpls>if>label-map

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

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

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

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

2.4.8.5 Interface Configuration for MPLS-TP

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

RFC 7213 provides guidelines for the usage of various Layer 2 next-hop resolution mechanisms with MPLS-TP. If protocols such as ARP are supported, then they should be used. However, in the case where no dynamic next hop resolution protocol is used, it should be possible to configure a unicast, multicast or broadcast next-hop MAC address. The rationale is to minimize the amount of configuration required for upstream nodes when downstream interfaces are changes. A default multicast MAC address for use by MPLS-TP point-to-point LSPs has been assigned by IANA (Value: 01-00-5e-90-00-00). This value is configurable on the router to support interoperability with third-party implementations that do not default to this value, and this no default value is implemented on the router.

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

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

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

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

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

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

MPLS-TP is only supported over Ethernet ports in Release 11.0. The system will block the association of an MPLS-TP LSP to an interface whose port is non-Ethernet.

If required, the IF_Num is configured under a MEP context under the MPLS interface. The **mpls-tp-mep** context is created under the interface as shown below. The *if-num* parameter, when concatenated with the Node ID, forms the IF_ID (as per RFC 6370), which is the identifier of this MEP. It is possible to configure this context whether the interface is IP numbered, IP unnumbered or mpls-tp unnumbered:

```
config
  router
    mpls
      interface <ip-int-name>
        mpls-tp-mep
          [no] ais-enable
          [no] if-num <if-num>
          [no] if-num-validation [enable | disable]
          ...
    exit
```

The **if-num-validation** command is used to enable or disable validation of the if-num in LSP Trace packet against the locally configured if-num for the interface over which the LSP Trace packet was received at the egress LER. This is because some implementations, do not perform interface validation for unnumbered MPLS-TP interfaces and instead set the if-num in the dsmap TLV to 0. The default is enabled.

AIS insertion is configured using the **ais-enable** command under the **mpls-tp-mep** context on an MPLS interface.

2.4.8.6 LER Configuration for MPLS-TP

2.4.8.6.1 LSP and Path Configuration

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

```
config
router
mpls
  lsp <xyz> [bypass-only | p2mp-lsp | mpls-tp <src-tunnel-num>]
    to node-id {<a.b.c.d> | <1.. .4,294,967,295>}
    dest-global-id <global-id>
    dest-tunnel-number <tunnel-num>
    [no] working-tp-path
      lsp-num <lsp-num>
      in-label <in-label>
      out-label <out-label> out-link <if-name>
        [next-hop <ipv4-address>]
    [no] mep
      [no] bfd-enable [cc | cc-cv]
      [no] bfd-trap-suppression
      [no] oam-template <name>
      [no] shutdown
    exit
  [no] shutdown
  exit
[no] protect-tp-path
  lsp-num <lsp-num>
  in-label <in-label>
  out-label <out-label> out-link <if-name>
    [next-hop <ipv4-address> ]
  [no] mep
    [no] bfd-enable [cc | cc-cv]
    [no] bfd-trap-suppression
    [no] oam-template <name>
    [no] protection-template <name>
    [no] shutdown
  exit
[no] shutdown
exit
```


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

<src-tunnel-num> is a mandatory create time parameter for mpls-tp tunnels, and has to be assigned by the user based on the configured range of tunnel ids. The *src-global-id* used for the LSP ID is derived from the node-wide *global-id* value configured under `config>router>mpls>mpls-tp`. A tunnel can not be brought up unless the *global-id* is configured.

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

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

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

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

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

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

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

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

There is no check made that a valid ARP entry exists before allowing a path to be un shut. Therefore, a path will only be held down if BFD is down. If static ARP is not configured for the interface, then it is assumed that dynamic ARP is used. The result is that if BFD is not configured, a path can come up before ARP resolution has completed for an interface. If BFD is not used, then it is recommended that the connectivity of the path is explicitly checked using on-demand CC/CV prior to sending user traffic on it.

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

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

- It is not mandatory to configure BFD or protection on an MPLS-TP path in order to bring the LSP up.
- If **bfd-enable cc** is configured, then CC-only mode using ACh channel 0x07 is used. If **bfd-enable cc_v** is configured, then BFD CC packets use channel 0x22 and CV packets use channel 0x23.
- Under the MEP context, the **bfd-trap-suppression** command allows the reception of AIS packets on the path to suppress BFD Down traps if a BFD session goes down on that path.

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

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

2.4.8.6.2 Support for Downstream Mapping Information

In order to validate the downstream mapping for an LSP, a node sending a DSMAP TLV must include the incoming and (optionally) outgoing IF_Num values for the interfaces that it expects the LSP to transit. Additionally, it will include the out-label for the LSP in the Label TLV for the DSMAP in the echo request message.

The incoming and outgoing if-num values correspond to the incoming and outgoing interfaces transited by an LSP at the next hop LER and LSR are configured using the **dsmap** command, as follows:

```
config
router
mpls
lsp
working-tp-path
mep
dsmap <in-if-num>[:<out-if-num>]
```

```
config
router
mpls
lsp
protect-tp-path
mep
dsmap <in-if-num>[:<out-if-num>]
```

```
config
router
mpls
mpls-tp
transit-path
```

```

forward-path
  mip
    dsmmap <in-if-num>[:<out-if-num>]
    exit
reverse-path
  mip
    dsmmap <in-if-num>[:<out-if-num>]
    exit

```

A node sending a DSMAP TLV will include these **in-if-num** and **out-if-num** (if configured) values. Additionally, it will include the out-label for the LSP in the Label TLV for the DSMAP in the echo request message.

2.4.8.6.3 Proactive CC/CV (using BFD) Configuration

Generally applicable proactive OAM parameters are configured using templates.

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

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

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

The BFD Template is configured as follows:

```

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

```

The parameters are as follows:

- **transmit-interval** *transmit-interval* and the **rx receive-interval**: These are the transmit and receive timers for BFD packets. If the template is used for MPLS-TP, then these are the timers used by CC packets. Values are in ms: 10 ms to 100 000 ms, with 1ms granularity. Default 10ms for CPM3 or better, 1 sec for other hardware. For MPLS-TP CV packets, a transmit interval of 1 s is always used.
- **multiplier** *multiplier*: Integer 3 to 20. Default: 3. This parameter is ignored for MPLS-TP combined cc-v BFD sessions, and the default of 3 used, as per RFC6428.
- **echo-receive** *echo-interval*: Sets the minimum echo receive interval (in ms), for a session. Values: 100 ms to 100 000 ms. Default: 100. This parameter is not used by a BFD session for MPLS-TP.
- **type cpm-np**: This selects the CPM network processor as the local termination point for the BFD session. This is enabled by default.

If the BFD timer values as shown above are changed in a template, any BFD sessions on MEPs to which that template is bound will try to renegotiate their timers to the new values.



Caution: The BFD implementations in some MPLS-TP peer nodes may not be able handle renegotiation, as allowed by Section 3.7.1 of RFC6428, and may take the BFD session down. This can result in undesired behavior, such as an unexpected protection switching event. We recommend that users of the system exercise caution when modifying the BFD timer values after a BFD session is up.

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

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

```
config
router
  mpls
    mpls-tp
      [no] oam-template <name>
      [no] bfd-template <name>
      [no] hold-time-down <interval>
      [no] hold-time-up <interval>
    exit
```

- **hold-time-down** *interval*: 0-5000 deciseconds, 10ms steps, default 0. This is equivalent to the standardized hold-off timer.

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

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

2.4.8.6.4 Protection templates and Linear Protection Configuration

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

A template is configured as follows:

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

The allowed values are as follows:

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

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

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

To minimize outage times, users should use the “mpls-tp protection command” (e.g. force/manual) to switch all the relevant MPLS-TP paths before executing the following commands:

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

2.4.8.7 Intermediate LSR Configuration for MPLS-TP LSPs

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

```
config
  router
    mpls
      mpls-tp
        transit-path <path-name>
          [no] path-id {lsp-num <lsp-num> | working-path | protect-path}
          [src-global-id <global-id>]
          src-node-id {<ipv4address> | <1.. .4,294,967,295>}
          src-tunnel-num <tunnel-num>
          [dest-global-id <global-id>]
          dest-node-id {<ipv4address> | <1.. .4,294,967,295>}
          [dest-tunnel-num <tunnel-num>]
        forward-path
          in-label <in-label> out-label <out-label>
          out-link <if-name> [next-hop <ipv4-next-hop>]
        reverse-path
          in-label <in-label> out-label <out-label>
          [out-link <if-name> [next-hop <ipv4-next-hop>]]
        [no] shutdown
```

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

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

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

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

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

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

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

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

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

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

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

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

2.4.9 MPLS-TP Show Commands

2.4.9.1 Static MPLS Labels

The following new commands show the details of the static MPLS labels.

show>router>mpls-labels>label <start-label> [<end-label> [in-use | <label-owner>]]

show>router>mpls-labels>label-range

An example output is as follows:

```
*A:mlstp-dutA# show router mpls
mpls          mpls-labels
*A:mlstp-dutA# show router mpls label
label          label-range
*A:7950 XRS-20# show router mpls-labels label-range
=====
Label Ranges
=====
Label Type      Start Label End Label   Aging      Available   Total
-----
Static          32          18431    -          18400      18400
Dynamic         18432       524287    0          505856     505856
```



```

      Seg-Route    0          0          -          0          505856
=====

```

2.4.9.2 MPLS-TP Tunnel Configuration

These commands show the configuration of a given tunnel.

show>router>mpls>tp-lsp

A sample output is as follows:

```

*A:mlstp-dutA# show router mpls tp-lsp
- tp-lsp [<lsp-name>] [status {up | down}] [from <ip-address> | to <ip-address>]
  [detail]
- tp-lsp [<lsp-name>] path [protect | working] [detail]
- tp-lsp [<lsp-name>] protection

```

```

<lsp-name>          : [32 chars max] - accepts * as wildcard char
<path>              : keyword - Display LSP path information.
<protection>        : keyword - Display LSP protection information.
<up | down>         : keywords - Specify state of the LSP
<ip-address>        : a.b.c.d
<detail>            : keyword - Display detailed information.

```

```

*A:mlstp-dutA# show router mpls tp-lsp
path
protection
to <a.b.c.d>
<lsp-name>
  "lsp-32" "lsp-33" "lsp-34" "lsp-35" "lsp-36" "lsp-37" "lsp-38" "lsp-39"
  "lsp-40" "lsp-41"
status {up | down}
from <ip-address>
detail

```

```

*A:mlstp-dutA# show router mpls tp-lsp "lsp-
"lsp-32" "lsp-33" "lsp-34" "lsp-35" "lsp-36" "lsp-37" "lsp-38" "lsp-39"
"lsp-40" "lsp-41"

```

```

*A:mlstp-dutA# show router mpls tp-lsp "lsp-32"

```

```

=====
MPLS MPLS-TP LSPs (Originating)
=====
LSP Name                                To                Tun    Protect  Adm  Opr
                                Id                Path
-----
lsp-32                                0.0.3.234         32     No       Up   Up
-----
LSPs : 1
=====
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" detail
=====

```

```

MPLS MPLS-TP LSPs (Originating) (Detail)
=====
Type : Originating
-----
LSP Name      : lsp-32
LSP Type      : MplsTp
From Node Id  : 0.0.3.233+
Adm State     : Up
LSP Up Time   : 0d 04:50:47
Transitions   : 1
DestGlobalId  : 42

LSP Tunnel ID : 32
To Node Id    : 0.0.3.234
Oper State    : Up
LSP Down Time : 0d 00:00:00
Path Changes   : 2
DestTunnelNum : 32

```

2.4.9.3 MPLS-TP Path configuration

This can reuse and augment the output of the current show commands for static LSPs. They should also show if BFD is enabled on a given path. If this referring to a transit path, this should also display (among others) the path-id (7 parameters) for a given transit-path-name, or the transit-path-name for a given the path-id (7 parameters)

show>router>mpls>tp-lsp>path

A sample output is as follows:

```

=====
*A:mlstp-dutA# show router mpls tp-lsp path
=====
MPLS-TP LSP Path Information
=====
LSP Name      : lsp-32
Admin State   : Up
To            : 0.0.3.234
Oper State    : Up

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              32        32        AtoB_1          Up    Down
Protect             2080      2080      AtoC_1          Up    Up
=====
LSP Name      : lsp-33
Admin State   : Up
To            : 0.0.3.234
Oper State    : Up

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              33        33        AtoB_1          Up    Down
Protect             2082      2082      AtoC_1          Up    Up
=====
LSP Name      : lsp-34
Admin State   : Up
To            : 0.0.3.234
Oper State    : Up

```

```

-----
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              34      34      AtoB_1          Up    Down
Protect             2084    2084    AtoC_1          Up    Up
=====
LSP Name      : lsp-35              To      : 0.0.3.234
Admin State   : Up                  Oper State : Up
-----

Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              35      35      AtoB_1          Up    Down
Protect             2086    2086    AtoC_1          Up    Up
=====
LSP Name      : lsp-36              To      : 0.0.3.234
Admin State   : Up                  Oper State : Up
-----

Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              36      36      AtoB_1          Up    Down
Protect             2088    2088    AtoC_1          Up    Up
=====
LSP Name      : lsp-37              To      : 0.0.3.234
Admin State   : Up                  Oper State : Up
-----

Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              37      37      AtoB_1          Up    Down
Protect             2090    2090    AtoC_1          Up    Up
=====
LSP Name      : lsp-38              To      : 0.0.3.234
Admin State   : Up                  Oper State : Up
-----

Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              38      38      AtoB_1          Up    Down
Protect             2092    2092    AtoC_1          Up    Up
=====
LSP Name      : lsp-39              To      : 0.0.3.234
Admin State   : Up                  Oper State : Up
-----

Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              39      39      AtoB_1          Up    Down
Protect             2094    2094    AtoC_1          Up    Up
=====
LSP Name      : lsp-40              To      : 0.0.3.234
Admin State   : Up                  Oper State : Up
-----

Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Working              40      40      AtoB_1          Up    Down
Protect             2096    2096    AtoC_1          Up    Up

```

```

=====
LSP Name       : lsp-41                               To       : 0.0.3.234
Admin State    : Up                                   Oper State : Up

-----
Path           NextHop           InLabel   OutLabel   Out I/F           Admin   Oper
-----
Working                41         41         AtoB_1         Up     Down
Protect              2098        2098        AtoC_1         Up     Up

*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path working

=====
MPLS-TP LSP Working Path Information
  LSP: "lsp-32"
=====
LSP Name       : lsp-32                               To       : 0.0.3.234
Admin State    : Up                                   Oper State : Up

-----
Path           NextHop           InLabel   OutLabel   Out I/F           Admin   Oper
-----
Working                32         32         AtoB_1         Up     Down

=====
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path protect

=====
MPLS-TP LSP Protect Path Information
  LSP: "lsp-32"
=====
LSP Name       : lsp-32                               To       : 0.0.3.234
Admin State    : Up                                   Oper State : Up

-----
Path           NextHop           InLabel   OutLabel   Out I/F           Admin   Oper
-----
Protect              2080        2080        AtoC_1         Up     Up

=====
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path protect detail

=====
MPLS-TP LSP Protect Path Information
  LSP: "lsp-32" (Detail)
=====
LSP Name       : lsp-32                               To       : 0.0.3.234
Admin State    : Up                                   Oper State : Up

Protect path information
-----
Path Type       : Protect                               LSP Num   : 2
Path Admin      : Up                                   Path Oper  : Up
Out Interface   : AtoC_1                               Next Hop Addr : n/a
In Label        : 2080                                   Out Label  : 2080
Path Up Time    : 0d 04:52:17                             Path Dn Time : 0d 00:00:00
Active Path     : Yes                                     Active Time  : 0d 00:52:56

MEP information
MEP State       : Up                                   BFD        : cc
OAM Templ       : privatebed-oam-template              CC Status   : inService

```

```
Protect Templ : privatebed-protection-template    CV Status      : unknown
RX PDU        : SF (1,1)                        WTR Count Down: 0 seconds
Defects       :                                TX PDU        : SF (1,1)
=====
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path working detail
=====
MPLS-TP LSP Working Path Information
  LSP: "lsp-32" (Detail)
=====
LSP Name      : lsp-32                          To              : 0.0.3.234
Admin State   : Up                              Oper State      : Up

Working path information
-----
Path Type     : Working                         LSP Num        : 1
Path Admin    : Up                             Path Oper      : Down
Down Reason   : ccFault ifDn
Out Interface : AtoB_1                         Next Hop Addr  : n/a
In Label      : 32                             Out Label      : 32
Path Up Time  : 0d 00:00:00                    Path Dn Time   : 0d 00:53:01
Active Path   : No                             Active Time    : n/a

MEP information
MEP State     : Up                             BFD            : cc
OAM Templ    : privatebed-oam-template        CC Status      : outOfService
                                                CV Status      : unknown
=====
*A:mlstp-dutA#
```

2.4.9.4 MPLS-TP Protection

These show the protection configuration for a given tunnel, which path in a tunnel is currently working and which is protect, and whether the working or protect is currently active.

show>router>mpls>tp-lsp>protection

A sample output is as follows:

```
*A:mlstp-dutA# show router mpls tp-lsp protection
=====
MPLS-TP LSP Protection Information
Legend: W-Working, P-Protect,
=====
LSP Name      Admin Oper Path  Ingr/Egr    Act. Rx PDU
              State State State Label      Path Tx PDU
-----
lsp-32        Up    Up    W Down    32/32       No  SF (1,1)
              P Up    2080/2080  Yes SF (1,1)
lsp-33        Up    Up    W Down    33/33       No  SF (1,1)
              P Up    2082/2082  Yes SF (1,1)
```

lsp-34	Up	Up	W Down	34/34	No	SF (1,1)
			P Up	2084/2084	Yes	SF (1,1)
lsp-35	Up	Up	W Down	35/35	No	SF (1,1)
			P Up	2086/2086	Yes	SF (1,1)
lsp-36	Up	Up	W Down	36/36	No	SF (1,1)
			P Up	2088/2088	Yes	SF (1,1)
lsp-37	Up	Up	W Down	37/37	No	SF (1,1)
			P Up	2090/2090	Yes	SF (1,1)
lsp-38	Up	Up	W Down	38/38	No	SF (1,1)
			P Up	2092/2092	Yes	SF (1,1)
lsp-39	Up	Up	W Down	39/39	No	SF (1,1)
			P Up	2094/2094	Yes	SF (1,1)
lsp-40	Up	Up	W Down	40/40	No	SF (1,1)
			P Up	2096/2096	Yes	SF (1,1)
lsp-41	Up	Up	W Down	41/41	No	SF (1,1)
			P Up	2098/2098	Yes	SF (1,1)

No. of MPLS-TP LSPs: 10
=====

2.4.9.5 MPLS TP Node Configuration

Displays the Global ID, Node ID and other general MPLS-TP configurations for the node.

show>router>mpls>mpls-tp

A sample output is as follows:

```
*A:mlstp-dutA# show router mpls mpls-tp
- mpls-tp
```

```
oam-template      - Display MPLS-TP OAM Template information
protection-tem*   - Display MPLS-TP Protection Template information
status            - Display MPLS-TP system configuration
transit-path      - Display MPLS-TP Tunnel information
```

```
*A:mlstp-dutA# show router mpls mpls-tp oam-template
```

```
=====
MPLS-TP OAM Templates
=====
```

```
Template Name : privatebed-oam-template Router ID      : 1
BFD Template  : privatebed-bfd-template Hold-Down Time: 0 centiseconds
                                           Hold-Up Time  : 20 deciseconds
```

```
*A:mlstp-dutA# show router mpls mpls-tp protection-template
```

```
=====
MPLS-TP Protection Templates
=====
```

```
Template Name : privatebed-protection-template Router ID      : 1
```

```

Protection Mode: one2one                      Direction      : bidirectional
Revertive      : revertive                    Wait-to-Restore: 300sec
Rapid-PSC-Timer: 10ms                        Slow-PSC-Timer  : 5sec
=====
*A:mlstp-dutA# show router mpls mpls-tp status

=====
MPLS-TP Status
=====
Admin Status   : Up
Global ID      : 42                      Node ID        : 0.0.3.233
Tunnel Id Min  : 1                      Tunnel Id Max   : 4096
=====
*A:mlstp-dutA# show router mpls mpls-tp transit-path
- transit-path [<path-name>] [detail]

<path-name>      : [32 chars max]
<detail>         : keyword - Display detailed information.


A:mlstp-dutC# show router mpls mpls-tp transit-path
- transit-path [<path-name>] [detail]

<path-name>      : [32 chars max]
<detail>         : keyword - Display detailed information.


A:mlstp-dutC# show router mpls mpls-tp transit-path
<path-name>
"tp-32"  "tp-33"  "tp-34"  "tp-35"  "tp-36"  "tp-37"  "tp-38"  "tp-39"
"tp-40"  "tp-41"
detail

A:mlstp-dutC# show router mpls mpls-tp transit-path "tp-32"

=====
MPLS-TP Transit tp-32 Path Information
=====
Path Name       : tp-32
Admin State     : Up                               Oper State      : Up

-----
Path           NextHop           InLabel   OutLabel   Out I/F
-----
FP              2080             2081      CtoB_1
RP              2081             2080      CtoA_1
=====
A:mlstp-dutC# show router mpls mpls-tp transit-path "tp-32" detail

=====
MPLS-TP Transit tp-32 Path Information (Detail)
=====
Path Name       : tp-32
Admin State     : Up                               Oper State      : Up
-----
Path ID configuration

```

```

Src Global ID : 42                      Dst Global ID : 42
Src Node ID   : 0.0.3.234              Dst Node ID   : 0.0.3.233
LSP Number    : 2                      Dst Tunnel Num: 32

Forward Path configuration
In Label      : 2080                    Out Label     : 2081
Out Interface : CtoB_1                  Next Hop Addr : n/a

Reverse Path configuration
In Label      : 2081                    Out Label     : 2080
Out Interface : CtoA_1                  Next Hop Addr : n/a
=====
A:mlstp-dutC#

```

2.4.9.6 MPLS-TP Interfaces

The existing `show>router>interface` command should be enhanced to display mpls-tp specific information.

The following is a sample output:

```

*A:mlstp-dutA# show router interface "AtoB_1"

=====
Interface Table (Router: Base)
=====
Interface-Name      Adm      Opr (v4/v6)  Mode      Port/SapId
  IP-Address                               PfxState
-----
AtoB_1              Down      Down/--      Network  1/2/3:1
  Unnumbered If[system]                               n/a
-----
Interfaces : 1

```

2.4.10 MPLS-TP Debug Commands

The following command provides the debug command for an MPLS-TP tunnel:

tools>dump>router>mpls>tp-tunnel <lsp-name> [clear]

The following is a sample output:

```

A:mlstp-dutA# tools dump router mpls tp-tunnel
- tp-tunnel <lsp-name> [clear]
- tp-tunnel id <tunnel-id> [clear]

<lsp-name>          : [32 chars max]
<tunnel-id>         : [1..61440]

```



```
<clear> : keyword - clear stats after reading

*A:mlstp-dutA# tools dump router mpls tp-tunnel "lsp-
"lsp-32" "lsp-33" "lsp-34" "lsp-35" "lsp-36" "lsp-37" "lsp-38" "lsp-39"
"lsp-40" "lsp-41"
*A:mlstp-dutA# tools dump router mpls tp-tunnel "lsp-32"

Idx: 1-32 (Up/Up): pgId 4, paths 2, operChg 1, Active: Protect
TunnelId: 42::0.0.3.233::32-42::0.0.3.234::32
PgState: Dn, Cnt/Tm: Dn 1/000 04:00:48.160 Up:3/000 00:01:25.840
MplsMsg: tpDn 0/000 00:00:00.000, tunDn 0/000 00:00:00.000
        wpDn 0/000 00:00:00.000, ppDn 0/000 00:00:00.000
        wpDel 0/000 00:00:00.000, ppDel 0/000 00:00:00.000
        tunUp 1/000 00:00:02.070

Paths:
Work (Up/Dn): Lsp 1, Lbl 32/32, If 2/128 (1/2/3 : 0.0.0.0)
  Tmpl: ptc: , oam: privatebed-oam-template (bfd: privatebed-bfd-template(np)-
10 ms)
  Bfd: Mode CC state Dn/Up handle 160005/0
  Bfd-CC (Cnt/Tm): Dn 1/000 04:00:48.160 Up:1/000 00:01:23.970
  State: Admin Up (1::1::1) port Up , if Dn , operChg 2
  DnReasons: ccFault ifDn

Protect (Up/Up): Lsp 2, Lbl 2080/2080, If 3/127 (5/1/1 : 0.0.0.0)
  Tmpl: ptc: privatebed-protection-template, oam: privatebed-oam-template (bfd:
privatebed-bfd-template(np)-10 ms)
  Bfd: Mode CC state Up/Up handle 160006/0
  Bfd-CC (Cnt/Tm): Dn 0/000 00:00:00.000 Up:1/000 00:01:25.410
  State: Admin Up (1::1::1) port Up , if Up , operChg 1

Aps: Rx - 5, raw 3616, nok 0(), txRaw - 3636, revert Y
Pdu: Rx - 0x1a-21::0101 (SF), Tx - 0x1a-21::0101 (SF)
State: PF:W:L LastEvt pdu (L-SFw/R-SFw)
Tmrs: slow
Defects: None Now: 000 05:02:19.130

Seq  Event  state  TxPdu  RxPdu  Dir  Act  Time
===  =====  =====  =====  =====  =====  =====  =====
000  start  UA:P:L  SF (0,0)  NR (0,0)  Tx-->  Work  000 00:00:02.080
001  pdu  UA:P:L  SF (0,0)  SF (0,0)  Rx<--  Work  000 00:01:24.860
002  pdu  UA:P:L  SF (0,0)  NR (0,0)  Rx<--  Work  000 00:01:26.860
003  pUp  NR  NR (0,0)  NR (0,0)  Tx-->  Work  000 00:01:27.440
004  pdu  NR  NR (0,0)  NR (0,0)  Rx<--  Work  000 00:01:28.760
005  wDn  PF:W:L  SF (1,1)  NR (0,0)  Tx-->  Prot  000 04:00:48.160
006  pdu  PF:W:L  SF (1,1)  NR (0,1)  Rx<--  Prot  000 04:00:48.160
007  pdu  PF:W:L  SF (1,1)  SF (1,1)  Rx<--  Prot  000 04:00:51.080
```

The following command shows the free mpls tunnel IDs available between two values, start-range and end-range.

```
tools>dump>router>mpls>free-tunnel-id <start-range> <end-range>
```

The following command provides a debug tool to view control-channel-status signaling packets.

```
*A:bksim1611# /debug service id 700 sdp 200:700 event-type ?{config-change |
oper-status-change | neighbor-discovery | control-channel-status}
```

```

*A:bksim1611# /debug service id 700 sdp 200:700 event-type control-channel-status

*A:bksim1611#
1 2012/08/31 09:56:12.09 EST MINOR: DEBUG #2001 Base PW STATUS SIG PKT (RX):
"PW STATUS SIG PKT (RX)::
Sdp Bind 200:700 Instance 3
  Version      : 0x0
  PW OAM Msg Type : 0x27
  Refresh Time  : 0xa
  Total TLV Length : 0x8
  Flags        : 0x0
  TLV Type     : 0x96a
  TLV Len      : 0x4
  PW Status Bits : 0x0
"

2 2012/08/31 09:56:22.09 EST MINOR: DEBUG #2001 Base PW STATUS SIG PKT (RX):
"PW STATUS SIG PKT (RX)::
Sdp Bind 200:700 Instance 3
  Version      : 0x0
  PW OAM Msg Type : 0x27
  Refresh Time  : 0xa
  Total TLV Length : 0x8
  Flags        : 0x0
  TLV Type     : 0x96a
  TLV Len      : 0x4
  PW Status Bits : 0x0
"

3 2012/08/31 09:56:29.44 EST MINOR: DEBUG #2001 Base PW STATUS SIG PKT (TX):
"PW STATUS SIG PKT (TX)::
Sdp Bind 200:700 Instance 3
  Version      : 0x0
  PW OAM Msg Type : 0x27
  Refresh Time  : 0x1e
  Total TLV Length : 0x8
  Flags        : 0x0
  TLV Type     : 0x96a
  TLV Len      : 0x4
  PW Status Bits : 0x0

```

2.5 Traffic Engineering

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

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

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

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

2.5.1 TE Metric (IS-IS and OSPF)

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

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

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

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

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

2.5.2 Admin Group Support on Facility Bypass Backup LSP

This feature provides for the inclusion of the LSP primary path admin-group constraints in the computation of a Fast ReRoute (FRR) facility bypass backup LSP to protect the primary LSP path by all nodes in the LSP path.

This feature is supported with the following LSP types and in both intra-area and inter-area TE where applicable:

- Primary path of a RSVP P2P LSP.
- S2L path of an RSVP P2MP LSP instance
- LSP template for an S2L path of an RSVP P2MP LSP instance.
- LSP template for auto-created RSVP P2P LSP in intra-area TE.

2.5.2.1 Procedures at Head-End Node

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

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

When this command is enabled at the ingress LER, the admin-group constraints configured in the context of the P2P LSP primary path, or the ones 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 thus propagates these constraints to the downstream nodes during the signaling of the LSP to allow them to include the admin-group constraints in the selection of the FRR backup LSP for protecting the LSP primary path.

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

The same admin-group constraints can be copied into the Session Attribute object. 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 regardless if the hop is strict or loose. These are governed strictly by the command:

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

In other words, the user may decide to copy the primary path admin-group constraints into the FAST_REROUTE object only, or into the Session Attribute object only, or into both.

The PLR rules for processing the admin-group constraints can make use of either of the two object admin-group constraints.

2.5.2.2 Procedures at PLR Node

The user enables the use of the 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 using the following global command:

config>router>mpls>admin-group-frr

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, then the PLR will read 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, then it just uses the admin-group constraint from the LSP and/or path level configurations.

Whether the PLR node is also the ingress LER or just an LSR for the protected LSP primary path, the outcome of the ingress LER configuration dictates the behavior of the PLR node and is summarized in [Table 6](#).

Table 6 Bypass LSP Admin-Group Constraint Behavior

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

Table 6 Bypass LSP Admin-Group Constraint Behavior (Continued)

	Ingress LER Configuration	Session Attribute	FRR Object	Bypass LSP at PLR (LER/LSF) follows admin-group constraints
2	frr-object lsp>propagate-admin-group lsp>frr>propagate-admin group	Admin color constraints sent	Admin color constraints sent	yes
3	frr-object lsp>propagate-admin group lsp>frr>no propagate-admin-group	Admin color constraints sent	Admin color constraints not sent	no
4	No frr-object lsp>propagate-admin group lsp>frr>propagate-admin-group	Admin color constraints sent	Not present	yes
5	No frr-object lsp>no propagate-admin group lsp>frr>propagate-admin-group	Admin color constraints not sent	Not present	no
6	No frr-object lsp>propagate-admin group lsp>frr>no propagate-admin-group	Admin color constraints sent	Not present	yes

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 among those that are already in use.

If none of the manual or dynamic bypass LSP 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 that includes or excludes the specified admin-group IDs.

If the user changes the configuration of the above command, it will not have any effect on existing bypass associations. The change will only apply to new attempts to find a valid bypass.

2.5.3 Diff-Serv Traffic Engineering

Diff-Serv traffic engineering provides the ability to manage bandwidth on a per Traffic Engineering (TE) class basis as per RFC 4124. In the base traffic engineering, LER computes LSP paths based on available BW of links on the path. Diff-Serv TE adds ability to perform this on a per TE class basis.

A TE class is a combination of Class Type and LSP priority. A Class Type is mapped to one or more system Forwarding Classes using a configuration profile. The operator sets different limits for admission control of LSPs in each TE class over each TE link. Eight TE classes are supported. Admission control of LSP paths bandwidth reservation is performed using the Maximum Allocation Bandwidth Constraint Model as per RFC 4125.

2.5.3.1 Mapping of Traffic to a Diff-Serv LSP

An LER allows the operator to map traffic to a Diff-Serv LSP using one of the following methods:

1. Explicit RSVP SDP configuration of a VLL, VPLS, or VPRN service
2. Class-based forwarding in an RSVP SDP. The operator can enable the checking by RSVP that a Forwarding Class (FC) mapping to an LSP under the SDP configuration is compatible with the Diff-Serv Class Type (CT) configuration for this LSP.
3. The **auto-bind-tunnel** RSVP-TE option in a VPRN service
4. Static routes with indirect next-hop being an RSVP LSP name

2.5.3.2 Admission Control of Classes

There are a couple of admission control decisions made when an LSP with a specified bandwidth is to be signaled. The first is in the head-end node. CSPF will only consider network links that have sufficient bandwidth. Link bandwidth information is provided by IGP TE advertisement by all nodes in that network.

Another decision made is local CAC and is performed when the RESV message for the LSP path is received in the reverse direction by a SR OS in that path. The bandwidth value selected by the egress LER will be checked against link bandwidth, otherwise the reservation is rejected. If accepted, the new value for the remaining link bandwidth will be advertised by IGP at the next advertisement event.

Both of these admission decisions are enhanced to be performed at the TE class level when Diff-Serv TE is enabled. In other words, CSPF in the head-end node will need to check the LSP bandwidth against the 'unreserved bandwidth' advertised for all links in the path of the LSP for that TE class which consists of a combination of a CT and a priority. Same for the admission control at SR OS receiving the Resv message.

2.5.3.2.1 Maximum Allocation Model

The admission control rules for this model are described in RFC 4125. Each CT shares a percentage of the Maximum Reservable Link Bandwidth through the user-configured BC for this CT. The Maximum Reservable Link Bandwidth is the link bandwidth multiplied by the RSVP interface subscription factor.

The sum of all BC values across all CTs will not exceed the Maximum Reservable Link Bandwidth. In other words, the following rule is enforced:

$$\text{SUM}(\text{BC}_c) \leq \text{Max-Reservable-Bandwidth}, 0 \leq c \leq 7$$

An LSP of class-type CT_c , setup priority p , holding priority h ($h \leq p$), and bandwidth B is admitted into a link if the following condition is satisfied:

$$B \leq \text{Unreserved Bandwidth for TE-Class}[i]$$

where TE-Class $[i]$ maps to $\langle \text{CT}_c, p \rangle$ in the definition of the TE classes on the node. The bandwidth reservation is effected at the holding priority, i.e., in TE-class $[j] = \langle \text{CT}_c, h \rangle$. Thus, the reserved bandwidth for CT_c and the unreserved bandwidth for the TE classes using CT_c are updated as follows:

$$\text{Reserved}(\text{CT}_c) = \text{Reserved}(\text{CT}_c) + B$$

$$\text{Unreserved TE-Class } [j] = \text{BC}_c - \text{SUM}(\text{Reserved}(\text{CT}_c, q)) \text{ for } 0 \leq q \leq h$$

$$\text{Unreserved TE-Class } [i] = \text{BC}_c - \text{SUM}(\text{Reserved}(\text{CT}_c, q)) \text{ for } 0 \leq q \leq p$$

The same is done to update the unreserved bandwidth for any other TE class making use of the same CT_c . These new values are advertised to the rest of the network at the next IGP-TE flooding.

When Diff-Serv is disabled on the node, this model degenerates into a single default CT internally with eight preemption priorities and a non-configurable BC equal to the Maximum Reservable Link Bandwidth. This would behave exactly like CT0 with eight preemption priorities and BC= Maximum Reservable Link Bandwidth if Diff-Serv was enabled.

2.5.3.2.2 Russian Doll Model

The RDM model is defined using the following equations:

$$\text{SUM (Reserved (CTc))} \leq \text{BCb},$$

where the SUM is across all values of **c** in the range **b** ≤ **c** ≤ (**MaxCT** - 1), and **BCb** is the bandwidth constraint of **CTb**.

BC0= **Max-Reservable-Bandwidth**, so that:

$$\text{SUM (Reserved(CTc))} \leq \text{Max-Reservable-Bandwidth},$$

where the **SUM** is across all values of **c** in the range 0 ≤ **c** ≤ (**MaxCT** - 1)

An LSP of class-type **CTc**, setup priority **p**, holding priority **h** (**h** ≤ **p**), and bandwidth **B** is admitted into a link if the following condition is satisfied:

$$\text{B} \leq \text{Unreserved Bandwidth for TE-Class [i]},$$

where **TE-Class [i]** maps to < **CTc**, **p** > in the definition of the TE classes on the node. The bandwidth reservation is effected at the holding priority, i.e., in **TE-class [j] = <CTc, h>**. Thus, the reserved bandwidth for CTc and the unreserved bandwidth for the TE classes using CTc are updated as follows:

$$\text{Reserved(CTc)} = \text{Reserved(CTc)} + \text{B}$$

$$\text{Unreserved TE-Class [j]} = \text{Unreserved (CTc, h)} = \text{Min [}$$

$$\text{BCc} - \text{SUM (Reserved (CTb, q) for } 0 \leq q \leq h, c \leq b \leq 7,$$

$$\text{BC(c-1)} - \text{SUM (Reserved (CTb, q) for } 0 \leq q \leq h, (c-1) \leq b \leq 7,$$

.....

$$\text{BC0} - \text{SUM (Reserved (CTb, q) for } 0 \leq q \leq h, 0 \leq b \leq 7]$$

$$\text{Unreserved TE-Class [i]} = \text{Unreserved (CTc, p)} = \text{Min [}$$

$$\text{BCc} - \text{SUM (Reserved (CTb, q) for } 0 \leq q \leq p, c \leq b \leq 7,$$

$$BC(c-1) - \text{SUM (Reserved (CT}_b, q) \text{ for } 0 \leq q \leq p, (c-1) \leq b \leq 7,$$

.....


$$BC0 - \text{SUM (Reserved (CT}_b, q) \text{ for } 0 \leq q \leq p, 0 \leq b \leq 7]$$

The same is done to update the unreserved bandwidth for any other TE class making use of the same CTc. These new values are advertised to the rest of the network at the next IGP-TE flooding.

Example CT Bandwidth Sharing with RDM

Below is a simple example with two CT values (CT0, CT1) and one priority 0 as shown in [Figure 25](#).

Figure 25 RDM with Two Class Types




Suppose CT1 bandwidth, or the CT1 percentage of Maximum Reservable Bandwidth to be more accurate is 100 Mb/s and CT2 bandwidth is 100 Mb/s and link bandwidth is 200 Mb/s. BC constraints can be calculated as follows.

$BC1 = \text{CT1 Bandwidth} = 100 \text{ Mb/s.}$

$BC0 = \{\text{CT1 Bandwidth}\} + \{\text{CT0 Bandwidth}\} = 200 \text{ Mb/s.}$

Suppose an LSP comes with CT1, setup and holding priorities of 0 and a bandwidth of 50 Mb/s.

Figure 26 First LSP Reservation



According to the RDM admission control policy:

$\text{Reserved}(\text{CT1}, 0) = 50 \leq 100 \text{ Mb/s}$

$\text{Reserved}(\text{CT0}, 0) + \text{Reserved}(\text{CT1}, 0) = 50 \leq 200 \text{ Mb/s}$

This results in the following unreserved bandwidth calculation.


$\text{Unreserved}(\text{CT1}, 0) = \text{BC1} - \text{Reserved}(\text{CT1}, 0) = 100 - 50 = 50 \text{ Mb/s}$

$\text{Unreserved}(\text{CT0}, 0) = \text{BC0} - \text{Reserved}(\text{CT0}, 0) - \text{Reserved}(\text{CT1}, 0) = 200 - 0 - 50 = 150 \text{ Mb/s}$.

The bandwidth reserved by a doll is not available to itself or any of the outer dolls.

Suppose now another LSP comes with CT0, setup and holding priorities of 0 and a bandwidth 120 Mb/s.

Figure 27 Second LSP Reservation



$\text{Reserved}(\text{CT0}, 0) = 120 \leq 150 \text{ Mb/s}$

$\text{Reserved (CT0, 0)} + \text{Reserved (CT1, 0)} = 120 + 50 = 170 \leq 200 \text{ Mb/s}$

$\text{Unreserved (CT0, 0)} = 150 - 120 = 30 \text{ Mb/s}$

If we simply checked BC1, the formula would yield the wrong results:

$\text{Unreserved (CT1, 0)} = \text{BC1} - \text{Reserved (CT1, 0)} = 100 - 50 = 50 \text{ Mb/s}$

Because of the encroaching of CT0 into CT1, we would need to deduct the overlapping reservation. This would then yield:

$\text{Unreserved (CT1, 0)} = \text{BC0} - \text{Reserved (CT0, 0)} - \text{Reserved (CT1, 0)} = 200 - 120 - 50 = 30 \text{ Mb/s}$, which is the correct figure.

Extending the formula with both equations:

$\text{Unreserved (CT1, 0)} = \text{Min} [\text{BC1} - \text{Reserved (CT1, 0)}, \text{BC0} - \text{Reserved (CT0, 0)} - \text{Reserved (CT1, 0)}] = \text{Min} [100 - 50, 200 - 120 - 50] = 30 \text{ Mb/s}$

An outer doll can encroach into an inner doll, reducing the bandwidth available for inner dolls.

2.5.3.3 RSVP Control Plane Extensions

RSVP will use the Class Type object to carry LSP class-type information during path setup. Eight values will be supported for class-types 0 through 7 as per RFC 4124. Class type 0 is the default class which is supported today on the router.

One or more forwarding classes will map to a Diff-Serv class type through a system level configuration.

2.5.3.4 IGP Extensions

IGP extensions are defined in RFC 4124. Diff-Serv TE advertises link available bandwidth, referred to as unreserved bandwidth, by OSPF TE or IS-IS TE on a per TE class basis. A TE class is a combination of a class type and an LSP priority. In order to reduce the amount of per TE class flooding required in the network, the number of TE classes is set to eight. This means that eight class types can be supported with a single priority or four class types with two priorities, etc. In that case, the operator configures the desired class type on the LSP such that RSVP-TE can signal it in the class-type object in the path message.

IGP will continue to advertise the existing Maximum Reservable Link Bandwidth TE parameter to mean the maximum bandwidth that can be booked on a given interface by all classes. The value advertised is adjusted with the link subscription factor.

2.5.3.5 Diff-Serv TE Configuration and Operation

2.5.3.5.1 RSVP Protocol Level

The following are the configuration steps at the RSVP protocol level:

1. The operator enables Diff-Serv TE by executing the **diffserv-te** command in the **config>router>rsvp** context. When this command is enabled, IS-IS and OSPF will start advertising available bandwidth for each TE class configured under the **diffserv-te** node. The operator can disable Diff-Serv TE globally by using the no form of the command.
2. The enabling or disabling of Diff-Serv on the system requires that the RSVP and MPLS protocol be shutdown. The operator must execute the **no shutdown** command in each context once all parameters under both protocols are defined. When saved in the configuration file, the **no shutdown** command is automatically inserted under both protocols to make sure they come up after a node reboot.
3. IGP will advertise the available bandwidth in each TE class in the unreserved bandwidth TE parameter for that class for each RSVP interface in the system.
4. In addition, IGP will continue to advertise the existing Maximum Reservable Link Bandwidth TE parameter so the maximum bandwidth that can be booked on a given interface by all classes. The value advertised is adjusted with the link subscription factor configured in the **config>router>rsvp>if>subscription percentage** context.
5. The operator can overbook (underbook) the maximum reservable bandwidth of a given CT by overbooking (underbooking) the interface maximum reservable bandwidth by configuring the appropriate value for the **subscription percentage** parameter.
6. The **diffserv-te** command will only have effect if the operator has already enabled traffic engineering at the IS-IS and/or OSPF routing protocol levels:
config>router>isis>traffic-engineering
and/or:
config>router>ospf>traffic-engineering

7. The following Diff-Serv TE parameters are configured globally under the **diffserv-te** node. They apply to all RSVP interfaces on the system. Once configured, these parameters can only be changed after shutting down the MPLS and RSVP protocols:

a. Definition of TE classes, TE Class = {Class Type (CT), LSP priority}. Eight TE classes can be supported. There is no default TE class once Diff-Serv is enabled. The operator must explicitly define each TE class. However, when Diff-Serv is disabled there will be an internal use of the default CT (CT0) and eight preemption priorities as shown in [Table 7](#).

Table 7 Internal TE Class Definition when Diff-Serv TE is Disabled

Class Type (CT internal)	LSP Priority
0	7
0	6
0	5
0	4
0	3
0	2
0	1
0	0

b. A mapping of the system forwarding class to CT. The default settings are shown in [Table 8](#).

Table 8 Default Mapping of Forwarding Class to TE Class

FC ID	FC Name	FC Designation	Class Type (CT)
7	Network Control	NC	7
6	High-1	H1	6
5	Expedited	EF	5
4	High-2	H2	4
3	Low-1	L1	3
2	Assured	AF	2
1	Low-2	L2	1

Table 8 Default Mapping of Forwarding Class to TE Class (Continued)

FC ID	FC Name	FC Designation	Class Type (CT)
0	Best Effort	BE	0

c. Configuration of the percentage of RSVP interface bandwidth each CT shares, for example, the Bandwidth Constraint (BC), using the **class-type-bw** command. The absolute value of the CT share of the interface bandwidth is derived as the percentage of the bandwidth advertised by IGP in the maximum reservable link bandwidth TE parameter, for example, the link bandwidth multiplied by the RSVP interface **subscription percentage** parameter. Note that this configuration also exists at the RSVP interface level and the interface specific configured value overrides the global configured value. The BC value can be changed at any time. The operator can specify the BC for a CT which is not used in any of the TE class definition but that does not get used by any LSP originating or transiting this node.

d. Configuration of the Admission Control Policy to be used: only the Maximum Allocation Model (MAM) is supported. The MAM value represents the bandwidth constraint models for the admission control of an LSP reservation to a link.

2.5.3.5.2 RSVP Interface Level

The following are the configuration steps at the RSVP interface level.

1. The operator configures the percentage of RSVP interface bandwidth each CT shares, for example, the BC, using the **class-type-bw** command. The value entered at the interface level overrides the global value configured under the **diffserv-te** node.
2. The operator can overbook (underbook) the maximum reservable bandwidth of a given CT by overbooking (underbooking) the interface maximum reservable bandwidth via configuring the appropriate value for the **subscription percentage** parameter in the **config>router>rsvp>interface** context.
3. Both the BC value and the subscription parameter can be changed at any time.

2.5.3.5.3 LSP and LSP Path Levels

The following are the configuration steps at the LSP and LSP path levels.

1. The operator configures the CT in which the LSP belongs by configuring the **class-type ct-number** command at the LSP level and/or the path level. The path level value overrides the LSP level value. By default, an LSP belongs to CT0.

2. Only one CT per LSP path is allowed per RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*. A multi-class LSP path is achieved through mapping multiple system Forwarding Classes to a CT.
3. The signaled CT of a dynamic bypass must always be CT0 regardless of the CT of the primary LSP path. The setup and hold priorities must be set to default values, for example, 7 and 0 respectively. This assumes that the operator configured a couple of TE classes, one which combines CT0 and a priority of 7 and the other which combines CT0 and a priority of 0. If not, the bypass LSP will not be signaled and will go into the down state.
4. The operator cannot configure the CT, setup priority, and holding priority of a manual bypass. They are always signaled with CT0 and the default setup and holding priorities.
5. The signaled CT, setup priority and holding priority of a detour LSP matches those of the primary LSP path it is associated with.
6. The operator can also configure the setup and holding priorities for each LSP path.
7. An LSP which does not have the CT explicitly configured will behave like a CT0 LSP when Diff-Serv is enabled.

If the operator configured a combination of a CT and a setup priority and/or a combination of a CT and a holding priority for an LSP path that are not supported by the user-defined TE classes, the LSP path will be kept in a down state and error code will be shown within the show command output for the LSP path.

2.5.4 Diff-Serv TE LSP Class Type Change under Failure

An option to configure a main Class Type (CT) and a backup CT for the primary path of a Diff-Serv TE LSP is provided. The main CT is used under normal operating conditions, for example, when the LSP is established the first time and when it gets re-optimized due to timer based or manual re-signal. The backup CT is used when the LSP retries under failure.

The use of backup Class Type (CT) by an LSP is enabled by executing the **config>router>mpls>lsp>primary>backup-class-type** *ct-number* command at the LSP primary path level.

When this option is enabled, the LSP will use the CT configured using the following commands (whichever is inherited at the primary path level) as the main CT:

- config>router>mpls>lsp>class-type *ct-number*
- config>router>mpls>lsp>primary>class-type *ct-number*

The main CT is used at initial establishment and during a manual or a timer based re-signal Make-Before-Break (MBB) of the LSP primary path. The backup CT is used temporarily to signal the LSP primary path when it fails and goes into retry.

Note that any valid values may be entered for the backup CT and main CT, but they cannot be the same. No check is performed to make sure that the backup CT is a lower CT in Diff-Serv Russian-Doll Model (RDM) admission control context.

The secondary paths of the same LSP are always signaled using the main CT as in existing implementation.

2.5.4.1 LSP Primary Path Retry Procedures

This feature behaves according to the following procedures.

- When a LSP primary path retries due a failure, for example, it fails after being in the up state, or undergoes any type of MBB, MPLS will retry a new path for the LSP using the main CT. If the first attempt failed, the head-end node performs subsequent retries using the backup CT. This procedure must be followed regardless if the currently used CT by this path is the main or backup CT. This applies to both CSPF and non-CSPF LSPs.
- The triggers for using the backup CT after the first retry attempt are:
 - A local interface failure or a control plane failure (hello timeout, etc.).
 - Receipt of a PathErr message with a notification of a FRR protection becoming active downstream and/or receipt of a Resv message with a 'Local-Protection-In-Use' flag set. This invokes the FRR Global Revertive MBB.
 - Receipt of a PathErr message with error code=25 (Notify) and sub-code=7 (Local link maintenance required) or a sub-code=8 (Local node maintenance required). This invokes the TE Graceful Shutdown MBB. Note that in this case, only a single attempt is performed by MBB as in current implementation; only the main CT will be retried.
 - Receipt of a Resv refresh message with the 'Preemption pending' flag set or a PathErr message with error code=34 (Reroute) and a value=1 (Reroute request soft preemption). This invokes the soft preemption MBB.
 - Receipt of a ResvTear message.
 - A configuration change MBB.

- When an unmapped LSP primary path goes into retry, it uses the main CT until the number of retries reaches the value of the new main-ct-retry-limit parameter. If the path did not come up, it must start using the backup CT at that point in time. By default, this parameter is set to infinite value. The new main-ct-retry-limit parameter has no effect on an LSP primary path, which retries due to a failure event. This parameter is configured using the **main-ct-retry-limit** command in the **config>router>mpls>lsp** context. If the user entered a value of the **main-ct-retry-limit** parameter that is greater than the LSP retry-limit, the number of retries will still stop when the LSP primary path reaches the value of the LSP retry-limit. In other words, the meaning of the LSP retry-limit parameter is not changed and always represents the upper bound on the number of retries. The unmapped LSP primary path behavior applies to both CSPF and non-CSPF LSPs.
- An unmapped LSP primary path is a path that never received a Resv in response to the first path message sent. This can occur when performing a “shut/no-shut” on the LSP or LSP primary path or when the node reboots. An unmapped LSP primary path goes into retry if the retry timer expired or the head-end node received a PathErr message before the retry timer expired.
- When the **clear>router>mpls>lsp** command is executed, the retry behavior for this LSP is the same as in the case of an unmapped LSP.
- If the value of the parameter main-ct-retry-limit is changed, the new value will only be used at the next time the LSP path is put into a “no-shut” state.
- The following is the behavior when the user changes the main or backup CT:
 - If the user changes the LSP level CT, all paths of the LSP are torn down and re-signaled in a break-before-make fashion. Specifically, the LSP primary path will be torn down and re-signaled even if it is currently using the backup CT.
 - If the user changes the main CT of the LSP primary path, the path will be torn down and re-signaled even if it is currently using the backup CT.
 - If the user changes the backup CT of an LSP primary path when the backup CT is in use, the path is torn down and is re-signaled.
 - If the user changes the backup CT of an LSP primary path when the backup CT is not in use, no action is taken. If however, the path was in global Revertive, gshut, or soft preemption MBB, the MBB is restarted. This actually means the first attempt will be with the main CT and subsequent ones, if any, with the new value of the backup CT.
 - Consider the following priority of the various MBB types from highest to lowest: Delayed Retry, Preemption, Global Revertive, Configuration Change, and TE Graceful Shutdown. If an MBB request occurs while a higher priority MBB is in progress, the latter MBB will be restarted. This actually means the first attempt will be with the main CT and subsequent ones, if any, with the new value of the backup CT.

- If the least-fill option is enabled at the LSP level, then CSPF must use least-fill equal cost path selection when the main or backup CT is used on the primary path.
- When the re-signal timer expires, CSPF will try to find a path with the main CT. The head-end node must re-signal the LSP even if the new path found by CSPF is identical to the existing one since the idea is to restore the main CT for the primary path. If a path with main CT is not found, the LSP remains on its current primary path using the backup CT. This means that the LSP primary path with the backup CT may no longer be the most optimal one. Furthermore, if the least-fill option was enabled at the LSP level, CSPF will not check if there is a more optimal path, with the backup CT, according to the least-fill criterion and will thus raise no trap to indicate the LSP path is eligible for least-fill re-optimization.
- When the user performs a manual re-signal of the primary path, CSPF will try to find a path with the main CT. The head-end node must re-signal the LSP as in current implementation.
- If a CPM switchover occurs while an the LSP primary path was in retry using the main or backup CT, for example, was still in operationally down state, the path retry will be restarted with the main CT until it comes up. This is because the LSP path retry count is not synchronized between the active and standby CPMs until the path becomes up.
- When the user configured secondary standby and non-standby paths on the same LSP, the switchover behavior between primary and secondary is the same as in existing implementation.


This feature is not supported on a P2MP LSP.

2.5.4.2 Bandwidth Sharing Across Class Types

In order to allow different levels of booking of network links under normal operating conditions and under failure conditions, it is necessary to allow sharing of bandwidth across class types.

This feature introduces the Russian-Doll Model (RDM) Diff-Serv TE admission control policy described in RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*. This mode is enabled using the following command: **config>router>rsvp>diffserv-te rdm**.

The Russian Doll Model (RDM) LSP admission control policy allows bandwidth sharing across Class Types (CTs). It provides a hierarchical model by which the reserved bandwidth of a CT is the sum of the reserved bandwidths of the numerically equal and higher CTs. [Figure 28](#) shows an example.

Figure 28 RDM Admission Control Policy Example

CT2 has a bandwidth constraint BC2 which represents a percentage of the maximum reservable link bandwidth. Both CT2 and CT1 can share BC1 which is the sum of the percentage of the maximum reservable bandwidth values configured for CT2 and CT1 respectively. Finally, CT2, CT1, and CT0 together can share BC0 which is the sum of the percentage of the maximum reservable bandwidth values configured for CT2, CT1, and CT0 respectively. The maximum value for BC0 is of course the maximum reservable link bandwidth.


What this means in practice is that CT0 LSPs can use up to BC0 in the absence of LSPs in CT1 and CT2. When this occurs and a CT2 LSP with a reservation less than or equal to BC2 requests admission, it is only admitted by preempting one or more CT0 LSPs of lower holding priority than this LSP setup priority. Otherwise, the reservation request for the CT2 LSP will be rejected.

It is required that multiple paths of the same LSP share common link bandwidth since they are signaled using the Shared Explicit (SE) style. Specifically, two instances of a primary path, one with the main CT and the other with the backup CT, must temporarily share bandwidth while MBB is in progress. Also, a primary path and one or many secondary paths of the same LSP must share bandwidth whether they are configured with the same or different CTs.

2.5.4.3 Downgrading the CT of Bandwidth Sharing LSP Paths

Consider a link configured with two class types CT0 and CT1 and making use of the RDM admission control model as shown in [Figure 29](#).

Figure 29 Sharing bandwidth when an LSP primary path is downgraded to backup CT




al_0210

Consider an LSP path Z occupying bandwidth B at CT1. $BC0$ being the sum of all CTs below it, the bandwidth occupied in CT1 is guaranteed to be available in CT0. Thus when new path X of the same LSP for CT0 is setup, it will use the same bandwidth B as used by path Z as shown in Figure 29 (a). When path Z is torn down the same bandwidth now occupies CT0 as shown in Figure 29 (b). Even if there were no new BW available in CT0 as can be seen in Figure 29 (c), path X can always share the bandwidth with path Z.

CSPF at the head-end node and CAC at the transit LSR node will share bandwidth of an existing path when its CT is downgraded in the new path of the same LSP.

2.5.4.4 Upgrading the CT of Bandwidth Sharing LSP Paths

When upgrading the CT the following issue can be apparent. Assume an LSP path X exists with CT0. An attempt is made to upgrade this path to a new path Z with CT1 using an MBB.

Figure 30 **Sharing Bandwidth When an LSP Primary Path is Upgraded to Main CT**

al_0211

In [Figure 30 \(a\)](#), if the path X occupies the bandwidth as shown it can not share the bandwidth with the new path Z being setup. If a condition exists, as shown in [Figure 30, \(b\)](#) the path Z can never be setup on this particular link.

Consider [Figure 30 \(c\)](#). The CT0 has a region that overlaps with CT1 as CT0 has incursion into CT1. This overlap can be shared. However, in order to find whether such an incursion has occurred and how large the region is, it is required to know the reserved bandwidths in each class. Currently, IGP-TE advertises only the unreserved bandwidths. Hence, it is not possible to compute these overlap regions at the head end during CSPF. Moreover, the head end needs to then try and mimic each of the traversed links exactly which increases the complexity.

CSPF at the head-end node will only attempt to signal the LSP path with an upgraded CT if the advertised bandwidth for that CT can accommodate the bandwidth. In other words, it will assume that in the worst case this path will not share bandwidth with another path of the same LSP using a lower CT.

2.6 Advanced MPLS/RSVP Features

- [Extending RSVP LSP to use Loopback Interfaces Other than Router-id](#)
- [LSP Path Change](#)
- [Manual LSP Path Switch](#)
- [Make-Before-Break \(MBB\) Procedures for LSP/Path Parameter Configuration Change](#)

- [Automatic Creation of RSVP-TE LSP Mesh](#)
- [RSVP-TE LSP Shortcut for IGP Resolution](#)
- [RSVP-TE LSP Signaling using LSP Template](#)
- [Shared Risk Link Groups](#)
- [TE Graceful Shutdown](#)
- [Soft Preemption of Diff-Serv RSVP LSP](#)
- [Least-Fill Bandwidth Rule in CSPF ECMP Selection](#)
- [Inter-Area TE LSP \(ERO Expansion Method\)](#)
- [Automatic Creation of a RSVP Mesh LSP](#)
- [Timer-based Reversion for RSVP-TE LSPs](#)
- [Automatic Creation of an RSVP One-Hop LSP](#)
- [MPLS Entropy Label](#)

2.6.1 Extending RSVP LSP to use Loopback Interfaces Other than Router-id

It is possible to configure the address of a loopback interface, other than the router-id, as the destination of an RSVP LSP, or a P2MP S2L sub-LSP. In the case of a CSPF LSP, CSPF searches for the best path that matches the constraints across all areas and 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 and levels of the IGP for that router-id; regardless of which area and level the router-id is reachable as an interface.

In addition, the user can now configure the address of a loopback interface, other than the router-id, 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 can 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, then CSPF will fail.

2.6.2 LSP Path Change

The **tools perform router mpls update-path {lsp lsp-name path current-path-name new-path new-path-name}** command instructs MPLS to replace the path of the primary or secondary LSP.

The primary or secondary LSP path is indirectly identified via the current-path-name value. In existing implementation, the same path name cannot be used more than once in a given LSP name.

This command is also supported on an SNMP interface.

This command applies to both CSPF LSP and to a non-CSPF LSP. However, it will only be honored when the specified current-path-name has the adaptive option enabled. The adaptive option can be enabled the LSP level or at the path level.

The new path must be first configured in CLI or provided via SNMP. The **configure router mpls path path-name** CLI command is used to enter the path.

The command fails if any of the following conditions are satisfied:

- The specified current-path-name of this LSP does not have the adaptive option enabled.
- The specified new-path-name value does not correspond to a previously defined path.
- The specified new-path-name value exists but is being used by any path of the same LSP, including this one.

When the command is executed, MPLS performs the following procedures:

- MPLS performs a single MBB attempt to move the LSP path to the new path.
- If the MBB is successful, MPLS updates the new path.
 - MPLS writes the corresponding NHLFE in the data path if this path is the current backup path for the primary.
 - If the current path is the active LSP path, it will update the path, write the new NHLFE in the data path, which will cause traffic to switch to the new path.
- If the MBB is not successful, the path retains its current value.
- The update-path MBB has the same priority as the manual re-signal MBB.

2.6.3 Manual LSP Path Switch

This feature provides a new command to move the path of an LSP from a standby secondary to another standby secondary.

The base version of the command allows the path of the LSP to move from a standby (or an active secondary) to another standby of the same priority. If a new standby path with a higher priority or a primary path comes up after the **tools perform** command is executed, the path re-evaluation command runs and the path is moved to the path specified by the outcome of the re-evaluation.

The CLI command for the base version is:

tools perform router mpls switch-path lsp *lsp-name* **path** *path-name*

The sticky version of the command can be used to move from a standby path to any other standby path regardless of priority. The LSP remains in the specified path until this path goes down or the user performs the no form of the **tools perform** command.

The CLI commands for the sticky version are:

tools perform router mpls force-switch-path lsp *lsp-name* **path** *path-name*

tools perform router mpls no force-switch-path lsp *lsp-name*

2.6.4 Make-Before-Break (MBB) Procedures for LSP/Path Parameter Configuration Change

When an LSP is switched from an existing working path to a new path, it is desirable to perform this in a hitless fashion. The Make-Before-Break (MBB) procedure consist of first signaling the new path when it is up, and having the ingress LER move the traffic to the new path. Only then the ingress LER tears down the original path.

MBB procedure is invoked during the following operations:

1. Timer based and manual re-signal of an LSP path.
2. Fast-ReRoute (FRR) global revertive procedures.
3. Soft Pre-emption of an LSP path.
4. Traffic-Engineering (TE) graceful shutdown procedures.
5. Update of secondary path due to an update to primary path SRLG.
6. LSP primary or secondary path name change.
7. LSP or path configuration parameter change.

In a prior implementation, item 7 covers the following parameters:

1. Changing the primary or secondary path **bandwidth** parameter on the fly.

2. Enabling the **frr** option for an LSP.

This feature extends the coverage of the MBB procedure to most of the other LSP level and Path level parameters as follows:

1. Changes to include/exclude of admin groups at LSP and path levels.
Enabling/disabling LSP level cspf option.
2. Enabling/disabling LSP level use-te-metric parameter when cspf option is enabled.
3. Enabling/disabling LSP level propagate-admin-group option.
4. Enabling/disabling LSP level hop-limit option in the fast-reroute context.
5. Enabling the LSP level least-fill option.
6. Enabling/disabling LSP level adspec option.
7. Changing between node-protect and “no node-protect” (link-protect) values in the LSP level fast-reroute option.
8. Changing LSP primary or secondary path priority values (setup-priority and hold-priority).
9. Changing LSP primary or secondary path class-type value and primary path backup-class-type value.
10. Changing LSP level and path level hop-limit parameter value.
11. Enabling/disabling primary or secondary path record and record-label options.

This feature is not supported on a manual bypass LSP.

P2MP Tree Level Make-before-break operation is supported if changes are made to the following parameters on LSP-Template:

- Changing Bandwidth on P2MP LSP-Template.
- Enabling Fast Re-Route on P2MP LSP-Template.

2.6.5 Automatic Creation of RSVP-TE LSP Mesh

This feature enables the automatic creation of an RSVP 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 auto-LSP of type mesh.

The user can associate multiple templates with the same or different peer prefix policies. Each application of an LSP template with a given prefix in the prefix list will result 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. Each instantiated LSP will have a unique LSP-id and a unique tunnel-ID.

Up to five (5) peer prefix policies can be associated with a given LSP template at all times. Each time the user executes the above command with the same or different prefix policy associations, or the user changes a prefix policy associated with an LSP template, the system re-evaluates the prefix policy. The outcome of the re-evaluation will tell MPLS if an existing LSP needs to be torn down or if a new LSP needs to be signaled to a destination address that is already in the TE database.

If a /32 prefix is added to (removed from) or if a prefix range is expanded (shrunk) in a prefix list associated with a LSP template, the same prefix policy re-evaluation described above is performed.

The trigger to signal the LSP is when the router with a router-id the matching a prefix in the prefix list appears in the Traffic Engineering database. The signaled LSP is installed in the Tunnel Table Manager (TTM) and is available to applications such as LDP-over-RSVP, resolution of BGP label routes, resolution of BGP, IGP, and static routes. It is, however, not available to be used as a provisioned SDP for explicit binding or auto-binding by services.

If the **one-hop** option is specified instead of a prefix policy, this command enables the automatic signaling of one-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. Although the provisioning model and CLI syntax differ from that of a mesh LSP only by the absence of a prefix list, the actual behavior is quite different. When the above command is executed, the TE database will keep track of each TE link that comes up 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. Thus, the **auto-lsp** command with the **one-hop** option will result in one or more LSPs signaled to the neighboring router.

An auto-created mesh or one-hop LSP can have egress statistics collected at the ingress LER by adding the **egress-statistics** node configuration into the LSP template. The user can also have ingress statistics collected at the egress LER using the same **ingress-statistics** node in CLI used with a provisioned LSP. The user must specify the full LSP name as signaled by the ingress LER in the RSVP session name field of the Session Attribute object in the received Path message.

2.6.6 RSVP-TE LSP Shortcut for IGP Resolution

RSVP-TE LSP shortcut for IGP route resolution allows forwarding of packets to IGP learned routes using an RSVP-TE LSP. This is also referred to as IGP shortcut. This feature is enabled by entering the following command at the IS-IS routing protocol level or at the OSPF routing protocol instance level:

- **config>router>isis>rsvp-shortcut**
- **config>router>ospf>rsvp-shortcut**

These commands instruct IS-IS or OSPF to include RSVP LSPs originating on this node and terminating on the router-id of a remote node as direct links with a metric equal to the operational metric provided by MPLS. Note that Dijkstra will always use the IGP metric to build the SPF tree and the LSP metric value does not update the SPF tree calculation. During the IP reach to determine the reachability of nodes and prefixes, LSPs are then overlaid and the LSP metric is used to determine the subset of paths which are equal lowest cost to reach a node or a prefix. If the user enabled the relative-metric option for this LSP, 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 computing the cost of a prefix which is resolved to the LSP.

When a prefix is resolved to a tunnel next-hop, the packet is sent labeled with the label stack corresponding to the NHLFE of the RSVP LSP. Any network event causing an RSVP LSP to go down will trigger a full SPF computation which may result in installing a new route over another RSVP LSP shortcut as tunnel next-hop or over a regular IP next-hop.

When rsvp-shortcut is enabled at the IGP instance level, all RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in **config>router>mpls>lsp>to**, corresponds to a router-id of a remote node. RSVP LSPs with a destination corresponding to an interface address or any other loopback interface address of a remote node are automatically not considered by IS-IS or OSPF. The user can, however, exclude a specific RSVP LSP from being used as a shortcut for resolving IGP routes by entering the command:

- **config>router>mpls>lsp>no igp-shortcut**

It is specifically recommended to disable igp-shortcut option on RSVP LSP which has the cspf option disabled unless the full explicit path of the LSP is provided in the path definition. MPLS tracks in RTM the destination or the first loose-hop in the path of a non CSPF LSP and as such this can cause bouncing when used within IGP shortcuts. The SPF in OSPF or IS-IS will only use RSVP LSPs as forwarding

adjacencies, IGP shortcuts, or as endpoints for LDP-over-RSVP. These applications of RSVP LSPs are mutually exclusive at the IGP instance level. If the user enabled two or more options in the same IGP instance, then forwarding adjacency takes precedence over the shortcut application that takes precedence over the LDP-over-RSVP application.

[Table 9](#) summarizes the outcome in terms of RSVP LSP role of mixing these configuration options.

Table 9 **RSVP LSP Role As Outcome of LSP level and IGP level configuration options**

	IGP Instance level configurations					
LSP level configuration	advertise-tunnel-link enabled / rsvp-shortcut enabled / ldp-over-rsvp enabled	advertise-tunnel-link enabled / rsvp-shortcut enabled / ldp-over-rsvp disabled	advertise-tunnel-link enabled / rsvp-shortcut disabled / ldp-over-rsvp disabled	advertise-tunnel-link disabled / rsvp-shortcut disabled / ldp-over-rsvp disabled	advertise-tunnel-link disabled / rsvp-shortcut enabled / ldp-over-rsvp enabled	advertise-tunnel-link disabled / rsvp-shortcut disabled / ldp-over-rsvp enabled
igp-shortcut enabled / ldp-over-rsvp enabled	Forwarding Adjacency	Forwarding Adjacency	Forwarding Adjacency	None	IGP Shortcut	LDP-over-RSVP
igp-shortcut enabled / ldp-over-rsvp disabled	Forwarding Adjacency	Forwarding Adjacency	Forwarding Adjacency	None	IGP Shortcut	None
igp-shortcut disabled / ldp-over-rsvp enabled	None	None	None	None	None	LDP-over-RSVP
igp-shortcut disabled / ldp-over-rsvp disabled	None	None	None	None	None	None

The resolution and forwarding of IPv6 prefixes to IPv4 IGP shortcuts is not supported.

The **no** form of this command disables the resolution of IGP routes using RSVP shortcuts.

2.6.6.1 Using LSP Relative Metric with IGP Shortcut

By default, the absolute metric of the LSP is used to compute the contribution of a IGP shortcut to the total cost of a prefix or a node after the SPF is complete. The absolute metric is the operational metric of the LSP populated by MPLS in the Tunnel Table Manager (TTM). This corresponds to the cumulative IGP-metric of the LSP path returned by CSPF or the static admin metric value of the LSP if the user configured one using the **config>router>mpls>lsp>metric** command. Note that MPLS populates the TTM with the maximum metric value of 16777215 in the case of a CSPF LSP using the TE-metric and a non-CSPF LSP with a loose or strict hop in the path. A non-CSPF LSP with an empty hop in the path definition returns the IGP cost for the destination of the LSP.

The user enables the use of the relative metric for an IGP shortcut with the following new CLI command:

```
config>router>mpls>lsp>igp-shortcut relative-metric [offset]
```

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 computing the cost of a prefix which is resolved to the LSP.

The offset value is optional and it defaults to zero. An offset value of zero is used when the **relative-metric** option is enabled without specifying the offset parameter value.

The minimum net cost for a prefix is capped to the value of one (1) after applying the offset:

$$\text{Prefix cost} = \max(1, \text{IGP cost} + \text{relative metric offset})$$

Note that the TTM continues to show the LSP operational metric as provided by MPLS. In other words, applications such as LDP-over-RSVP (when IGP shortcut is disabled) and BGP and static route shortcuts will continue to use the LSP operational metric.

The **relative-metric** option is mutually exclusive with the **lfa-protect** or the **lfa-only** options. In other words, an LSP with the **relative-metric** option enabled cannot be included in the LFA SPF and vice-versa when the **rsvp-shortcut** option is enabled in the IGP.

Finally, it should be noted that the **relative-metric** option is ignored when forwarding adjacency is enabled in IS-IS or OSPF by configuring the **advertise-tunnel-link** option. In this case, IGP advertises the LSP as a point-to-point unnumbered link along with the LSP operational metric capped to the maximum link metric allowed in that IGP.

The resolution and forwarding of IPv6 prefixes to IPv4 forwarding adjacency LSP is not supported.

2.6.6.2 ECMP Considerations

When ECMP is enabled on the system and multiple equal-cost paths exist for a prefix, the following selection criteria are used to pick up the set of next-hops to program in the data path:

- for a destination = tunnel-endpoint (including external prefixes with tunnel-endpoint as the next-hop):
 - select tunnel with lowest tunnel-index (ip next-hop is never used in this case)
- for a destination != tunnel-endpoint:
 - exclude LSPs with metric higher than underlying IGP cost between the endpoint of the LSP
 - prefer tunnel next-hop over ip next-hop
 - within tunnel next-hops:
 - i. select lowest endpoint to destination cost
 - ii. if same endpoint to destination cost, select lowest endpoint node router-id
 - iii. if same router-id, select lowest tunnel-index
 - within ip next-hops:
 - select lowest downstream router-id
 - if same downstream router-id, select lowest interface-index
- Note though no ECMP is performed across both the IP and tunnel next-hops the tunnel endpoint lies in one of the shortest IGP paths for that prefix. In that case, the tunnel next-hop is always selected as long as the prefix cost using the tunnel is equal or lower than the IGP cost.

The ingress IOM will spray the packets for a prefix over the set of tunnel next-hops and IP next-hops based on the hashing routine currently supported for IPv4 packets.

2.6.6.3 Handling of Control Packets

All control plane packets that require an RTM lookup and whose destination is reachable over the RSVP shortcut will be forwarded over the shortcut. This is because RTM keeps a single route entry for each prefix unless there is ECMP over different outgoing interfaces.

Interface bound control packets are not impacted by the RSVP shortcut since RSVP LSPs with a destination address different than the router-id are not included by IGP in its SPF calculation.

2.6.6.4 Forwarding Adjacency

The forwarding adjacency feature can be enabled independently from the IGP shortcut feature in CLI. To enable forwarding adjacency, the user enters the following command in IS-IS or OSPF:

- **config>router>isis>advertise-tunnel-link**
- **config>router>ospf>advertise-tunnel-link**

If both **rsvp-shortcut** and **advertise-tunnel-link** options are enabled for a given IGP instance, then the **advertise-tunnel-link** will win. With this feature, ISIS or OSPF advertises an RSVP LSP as a link so that other routers in the network can include it in their SPF computations. The RSVP LSP is advertised as an unnumbered point-to-point link and the link LSP/LSA has no Traffic Engineering opaque sub-TLVs as per RFC 3906 *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels*.

The forwarding adjacency feature can be enabled independently from the IGP shortcut feature in CLI. If both **rsvp-shortcut** and **advertise-tunnel-link** options are enabled for a given IGP instance, then the **advertise-tunnel-link** will win.

When the forwarding adjacency feature is enabled, each node advertises a p2p unnumbered link for each best metric tunnel to the router-id of any endpoint node. The node does not include the tunnels as IGP shortcuts in SPF computation directly. Instead, when the LSA/LSP advertising the corresponding P2P unnumbered link is installed in the local routing database, then the node performs an SPF using it like any other link LSA/LSP. The link bi-directional check requires that a link, regular link or tunnel link, exists in the reverse direction for the tunnel to be used in SPF.

Note that the **igp-shortcut** option under the LSP name governs the use of the LSP with both the **rsvp-shortcut** and the **advertise-tunnel-link** options in IGP. The interactions of these options are summarized in [Table 10](#):

Table 10 **Impact of LSP level configuration on IGP shortcut and forwarding adjacency features**

LSP level configuration	Actions with IGP Shortcut Feature	Actions with Forwarding Adjacency Feature
igp-shortcut	Tunnel is used in main SPF, but is not used in LFA SPF	Tunnel is advertised as p2p link if it has best LSP metric, is used in main SPF if advertised, but is not used in LFA SPF
igp-shortcut lfa-protect	Tunnel is used in main SPF, and is used in LFA SPF	Tunnel is advertised as p2p link if it has best LSP metric, is used in main SPF if advertised, and is used in LFA SPF regardless if it is advertised or not
igp-shortcut lfa-only	Tunnel is not used in main SPF, but is used in LFA SPF	Tunnel is not advertised as p2p link, if not used in main SPF, but is used in LFA SPF

2.6.6.5 LDP Forwarding over IGP Shortcut

The user can enable LDP FECs over IGP shortcuts by configuring T-LDP sessions to the destination of the RSVP LSP. In this case, LDP FEC is tunneled over the RSVP LSP, effectively implementing LDP-over-RSVP without having to enable the **ldp-over-rsvp** option in OSPF or IS-IS. The **ldp-over-rsvp** and **igp-shortcut** options are mutually exclusive under OSPF or IS-IS.

2.6.6.6 Handling of Multicast Packets

This feature supports multicast Reverse-Path Check (RPF) in the presence of IGP shortcuts. When the multicast source for a packet is reachable via an IGP shortcut, the RPF check fails since PIM requires a bi-directional path to the source but IGP shortcuts are unidirectional.

The implementation of the IGP shortcut feature provides IGP with the capability to populate the multicast RTM with the prefix IP next-hop when both the **rsvp-shortcut** option and the **multicast-import** option are enabled in IGP.

This change is made possible with the enhancement introduced by which SPF keeps track of both the direct first hop and the tunneled first hop of a node that is added to the Dijkstra tree.

Note that IGP will not pass LFA next-hop information to the mcast RTM in this case. Only ECMP next-hops are passed. As a consequence, features such as PIM Multicast-Only FRR (MoFRR) will only work with ECMP next-hops when IGP shortcuts are enabled.

Finally, note that the concurrent enabling of the **advertise-tunnel-link** option and the **multicast-import** option will result a multicast RTM that is a copy of the unicast RTM and is thus populated with mix of IP and tunnel NHs. RPF will succeed for a prefix resolved to a IP NH, but will fail for a prefix resolved to a tunnel NH. [Table 11](#) summarizes the interaction of the **rsvp-shortcut** and **advertise-tunnel-link** options with unicast and multicast RTMs.

Table 11 Impact of IGP Shortcut and Forwarding Adjacency on Unicast and Multicast RTM

		Unicast RTM (Primary SPF)	Multicast RTM (Primary SPF)	Unicast RTM (LFA SPF)	Multicast RTM (LFA SPF)
OSPF	rsvp-shortcut	√	√ ¹	√	X ³
	advertise-tunnel-link	√	√ ²	√	√ ⁴
IS-IS	rsvp-shortcut	√	√ ¹	√	X ³
	advertise-tunnel-link	√	√ ²	√	√ ⁴

Notes:

1. Multicast RTM is different from unicast RTM as it is populated with IP NHs only, including ECMP IP NHs. RPF check can be performed for all prefixes.
2. Multicast RTM is a copy of the unicast RTM and is thus populated with mix of IP and tunnel NHs. RPF will succeed for a prefix resolved to a IP NH but will fail for a prefix resolved to a tunnel NH.
3. LFA NH is not computed for the IP primary next-hop of a prefix passed to multicast RTM even if the same IP primary next-hop ends up being installed in the unicast RTM. The LFA next-hop will, however, be computed and installed in the unicast RTM for a primary IP next-hop of a prefix.
4. Multicast RTM is a copy of the unicast RTM and is thus populated with mix of IP and tunnel LFA NHs. RPF will succeed for a prefix resolved to a primary or LFA IP NH but will fail for a prefix resolved to a primary or LFA tunnel NH.

2.6.6.6.1 Disabling TTL Propagation in an LSP Shortcut

This feature provides the option for disabling TTL propagation from a transit or a locally generated IP packet header into the LSP label stack when an RSVP LSP is used as a shortcut for BGP next-hop resolution, a static-route next-hop resolution, or for an IGP route resolution.

A transit packet is a packet received from an IP interface and forwarded over the LSP shortcut at ingress LER.

A locally-generated IP packet is any control plane packet generated from the CPM and forwarded over the LSP shortcut at ingress LER.

TTL handling can be configured for all RSVP LSP shortcuts originating on an ingress LER using the following global commands:

```
config>router>mpls>[no] shortcut-transit-ttl-propagate  
config>router>mpls>[no] shortcut-local-ttl-propagate
```

These commands apply to all RSVP LSPs which are used to resolve static routes, BGP routes, and IGP routes.

When the **no** form of the above command is enabled for local packets, TTL propagation is disabled on all locally generated IP packets, including ICMP Ping, trace route, and OAM packets that are destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack. This is referred to as pipe mode.

Similarly, when the **no** form is enabled for transit packets, TTL propagation is disabled on all IP packets received on any IES interface and destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack.

2.6.7 RSVP-TE LSP Signaling using LSP Template

An LSP template can be used for signaling RSVP-TE LSP to far-end PE node that is detected based on auto-discovery method by a client application. RSVP-TE P2MP LSP signaling based on LSP template is supported for Multicast VPN application on SR OS platform. LSP template avoids an explicit LSP or LSP S2L configuration for a node that is dynamically added as a receiver.

LSP template has option to configure traffic engineering parameters that apply to LSP that is setup using the template. Traffic engineering options that are currently supported are:

- adaptive
- admin-group
- bandwidth
- CSPF calculation
- fast-reroute

- hop-limit
- record-label
- retry-timer

2.6.8 Shared Risk Link Groups

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

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

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

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

2.6.8.1 Enabling Disjoint Backup Paths

A typical application of the SRLG feature is to provide for an automatic placement of secondary backup LSPs or FRR bypass/detour LSPs that minimizes the probability of fate sharing with the path of the primary LSP ([Figure 31](#)).

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

- For primary/standby SRLG disjoint configuration:
 - Create an SRLG-group, similar to admin groups.
 - Link the SRLG-group to MPLS interfaces.
 - Configure primary and secondary LSP paths and enable SRLG on the secondary LSP path. Note that the SRLG secondary LSP path(s) will *always* perform a strict CSPF query. The **srlg-frr** command is irrelevant in this case (see [srlg-frr](#)).
- For FRR detours/bypass SRLG disjoint configuration:
 - Create an SRLG group, similar to admin groups.
 - Link the SRLG group to MPLS interfaces.
 - Enable the **srlg-frr** (strict/non-strict) option, which is a system-wide parameter, and it force every LSP path CSPF calculation, to take the configured SRLG membership(s) (and propagated through the IGP opaque-te-database) into account.
 - Configure primary FRR (one-to-one/facility) LSP path(s). Consider that each PLR will create a detour/bypass that will only avoid the SRLG membership(s) configured on the primary LSP path egress interface. In a one-to-one case, detour-detour merging is out of the control of the PLR, thus the latter will not ensure that its detour will be prohibited to merge with a colliding one. For facility bypass, with the presence of several bypass type to bind to, the following priority rules will be followed:
 1. Manual bypass disjoint
 2. Manual bypass non-disjoint (eligible only if srlg-frr is non-strict)
 3. Dynamic disjoint
 4. Dynamic non-disjoint (eligible only if srlg-frr is non-strict)Non-CSPF manual bypass is not considered.


Figure 31 Shared Risk Link Groups

Fig.33


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

2.6.8.2 SRLG Penalty Weights for Detour and Bypass LSPs

The likelihood of paths with links sharing SRLG values with a primary path being used by a bypass or detour LSP can be configured if a penalty weight is specified for the link. The higher the penalty weight, the less desirable it is to use the link with a given SRLG.

[Figure 32](#) illustrates the operation of SRLG penalty weights.

Figure 32 SRLG Penalty Weight Operation



24823

The primary LSP path includes a link between A and D with SRLG (1) and (2). The bypass around this link through nodes B and C includes links (a) and (d), which are members of SRLG (1), and links (b) and (c), which are members of SRLG 2. If the link metrics are equal, then this gives four ECMP paths from A to D via B and C:

- (a), (d), (e)
- (a), (c), (e)
- (b), (c), (e)
- (b), (d), (e)

Two of these paths include undesirable (from a reliability perspective) link (c). SRLG penalty weights or costs can be used to provide a tiebreaker between these paths so that the path including (c) is less likely to be chosen. For example, if the penalty associated with SRLG (1) is 5, and the penalty associated with SRLG (2) is 10, and the penalty associated with SRLG (3) is 1, then the cumulative penalty of each of the paths above is calculated by summing the penalty weights for each SRLG that a path has in common with the primary path:

- (a), (d), (e) = 10
- (a), (c), (e) = 15
- (b), (c), (e) = 20
- (b), (d), (e) = 15

Therefore path (a), (d), (e) is chosen since it has the lowest cumulative penalty.

Penalties are applied by summing the values for SRLGs in common with the protected part of the primary path.

A user can define a penalty weight value associate with an SRLG group using the **penalty-weight** parameter of the **srlg-group** command under the **configure>router-if-attribute** context. If an SRLG penalty weight is configured, then CSPF will include the SRLG penalty weight in the computation of an FRR detour or bypass for protecting the primary LSP path at a PLR node. Links with a higher SRLG penalty should be more likely to be pruned than links with a lower SRLG penalty.

Note that the configured penalty weight is not advertised in the IGP.

An SRLG penalty weight is applicable whenever an SRLG group is applied to an interface, including in the static SRLG database. However, penalty weights are used in bypass and detour path computation only when the **srlg-frr (loose)** flag is enabled.

2.6.8.3 Static Configurations of SRLG Memberships

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

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

Note, however, that the SRLG secondary path computation and FRR bypass/detour path computation remains unchanged.

There are deployments where the SR OS will interoperate with routers that do not implement the SRLG membership advertisement via IGP SRLG TLV or sub-TLV.

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

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

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

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

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

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

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

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

2.6.9 TE Graceful Shutdown

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

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

2.6.10 Soft Preemption of Diff-Serv RSVP LSP

A Diff-Serv LSP can preempt another LSP of the same or of a different CT if its setup priority is strictly higher (numerically lower) than the holding priority of that other LSP.

2.6.11 Least-Fill Bandwidth Rule in CSPF ECMP Selection

When multiples equal-cost paths satisfy the constraints of a given RSVP LSP path, CSPF in the router head-end node will select a path so that LSP bandwidth is balanced across the network links. In releases prior to R7.0, CSPF used a random number generator to select the path and returned it to MPLS. In the course of time, this method actually balances the number of LSP paths over the links in the network; it does not necessarily balance the bandwidth across those links.

The least-fill path selection algorithm identifies the single link in each of the equal cost paths which has the least available bandwidth in proportion to its maximum reserved bandwidth. It then selects the path which has the largest value of this figure. The net affect of this algorithm is that LSP paths will be spread over the network links over time such that percentage link utilization is balanced. When the least-fill option is enabled on an LSP, during a manual reset CSPF will apply this method to all path calculations of the LSP, also at the time of the initial configuration.

2.6.12 Inter-Area TE LSP (ERO Expansion Method)

Inter-area contiguous LSP scheme provides end-to-end TE path. Each transit node in an area can set up a TE path LSP based on TE information available within its local area.

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

Area border router on receiving a PATH message with loose hop ERO does a partial CSPF calculation to the next domain border router as loose hop or CSPF to reach the final destination.


2.6.12.1 Area Border Node FRR Protection for Inter-Area LSP

This feature enhances the prior implementation of an inter-area RSVP P2P LSP by making the ABR selection automatic at the ingress LER. The user will not need to include the ABR as a loose-hop in the LSP path definition.

CSPF adds a new capability to compute all segments of a multi-segment intra-area or inter-area LSP path in one operation. In previous releases, MPLS makes a request to CSPF for each segment separately.

Figure 33 illustrates the role of each node in the signaling of an inter-area LSP with automatic ABR node selection.

Figure 33 Automatic ABR Node Selection for Inter-Area LSP




CSPF for an inter-area LSP operates as follows:

1. CSPF in the Ingress LER node determines that an LSP is inter-area by doing a route lookup with the destination address of a P2P LSP (i.e., 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 as inter-area.
2. When the path of the LSP is empty, CPSF will compute a single-segment intra-area path to an ABR node that advertised a prefix matching with the destination address of the LSP.
3. When the path of the LSP contains one or more hops, CSPF will compute a multi-segment intra-area path including the hops that are in the area of the Ingress LER node.
4. When all hops are in the area of the ingress LER node, the calculated path ends on an ABR node that advertised a prefix matching with the destination address of the LSP.
5. When there are one or more hops that are not in the area of the ingress LER node, the calculated path ends on an ABR node that advertised a prefix matching with the first hop-address that is not in the area of the ingress LER node.
6. 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.

This behavior can be illustrated in the [Figure 34](#). The TE link between ABR nodes D and E is in area 0. When node C computes the path for LSP from C to B which path specified nodes C and D as loose hops, it would fail the path computation if CSPF attempted a path all the way to the last hop in the local area, node E. Instead, CSPF stops the path at node A which will further expand the ERO by including link D-E as part of the path in area 0.

Figure 34 CSPF for an Inter-area LSP



al_0905

7. If there is more than 1 ABR that advertised a prefix, CSPF will calculate a path for all ABRs. Only the shortest path will be withheld. If more than one path has the shortest path, CSPF will pick a path randomly or based on the least-fill criterion if enabled. If more than one ABR satisfies the least-fill criterion, CSPF will also pick one path randomly.
8. The path for an intra-area LSP path will not be able to exit and re-enter the local area of the ingress LER. This behavior was possible in prior implementation when the user specified a loose hop outside of the local area or when the only available path was via TE links outside of the local area.

2.6.12.1.1 Rerouting of Inter-Area LSP

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

One new behavior introduced by the automatic selection of ABR is the ability of the ingress LER to reroute an inter-area LSP primary path via a different ABR in the following situations:

- When the local exit ABR node fails, There are two cases to consider:
 - The primary path is not protected at the ABR and is thus torn down by the previous hop in the path. In this case the ingress LER will retry the LSP primary path via the ABR which currently has the best path for the destination prefix of the LSP.
 - The primary path is protected at the ABR with a manual or dynamic bypass LSP. In this case the ingress LER will receive a Path Error message with a notification of a protection becoming active downstream and a RESV with a *Local-Protection-In-Use* flag set. At the receipt of first of these two messages, the ingress LER will then perform a Global Revertive Make-Before-Break (MBB) to re-optimize the LSP primary path via the ABR which 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. In this case, the ingress LER will perform a MBB to re-optimize the LSP primary path via the ABR which 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 re-optimization of the LSP path in the case of the receipt of the IS-IS overload bit.

2.6.12.1.2 Behavior of MPLS Options in Inter-Area LSP

The automatic ABR selection for an inter-area LSP does not change prior implementation inter-area LSP behavior of many of the LSP and path level options. There is, however, a number of enhancements introduced by the automatic ABR selection feature as explained in the following.

- 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 and soft preemption features will continue 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 will not be able to avoid the link or node when it performs the MBB since it computes the path to the local ABR exit router only. There is, however, an exception to this for the TE graceful shutdown case only. An enhancement has been added to cause the upstream ABR nodes in the current path of the LSP to record the link or node to avoid and will use it in subsequent ERO expansions. This means that if the ingress LER computes a new MBB path which goes via the same exit ABR router as the current path and all ABR upstream nodes of the node or link which originated the PathErr message are also selected in the new MBB path when the ERO is expanded, the new path will indeed avoid this link or node. The latter is a new behavior introduced with the automatic ABR selection feature.
- The support of MBB to avoid the ABR node when the node is put into TE Graceful Shutdown is a new behavior introduced with the automatic ABR selection feature.
- The **use-te-metric** option in CSPF cannot be propagated across the area boundary and thus will operate within the scope of the local area of the ingress LER node. This is a new behavior introduced with the automatic ABR selection feature.
- The **srlg** option on bypass LSP will continue to operate locally at each PLR within each area. The PLR node protecting the ABR will check the SRLG constraint for the path of the bypass within the local area.
- The **srlg** option on secondary path is allowed to operate within the scope of the local area of the ingress LER node 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 node in the local area.
- The PLR node must indicate to CSPF that a request to 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.

- The **propagate-admin-group** option under the LSP will still need to be enabled on the inter-area LSP if the user wants to have admin-groups propagated across the areas.
- With the automatic ABR selection feature, timer based re-signal of the inter-area LSP path will be supported and will re-signal the path if the cost of the path segment to the local exit ABR changed. The cost shown for the inter-area LSP at ingress LER will be the cost of the path segments to the ABR node.

2.6.12.2 Inter-Area LSP support of OSPF Virtual Links

The OSPF virtual link extends area 0 for a router that is not connected to area 0. As a result, it makes all prefixes in area 0 reachable via an intra-area path but in reality, they are not 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 is directly connected to this node.

This inter-area LSP feature can identify when the destination of an LSP is reachable via a virtual link. In that case, CSPF will automatically compute and signal an inter-area LSP via the ABR nodes that is 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 will compute and signal an intra-area LSP using the transit area TE links, even when the destination router-id is only part of area 0.

2.6.12.3 Area Border Node FRR Protection for Inter-Area LSP

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

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

Dynamic bypass protection provides for the automatic computation, signaling, and association with the primary path of an inter-area P2P LSP to provide ABR node protection. [Figure 35](#) illustrates the role of each node in the ABR node protection using a dynamic bypass LSP.

Figure 35 ABR Node Protection Using Dynamic Bypass LSP

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

- The PLR node must inspect the node-id RRO 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 node and with the XRO set to exclude the node-id of the protected ABR node.
- The request to CSPF is for a path to the merge-point (i.e., the next-next-hop in the RRO received in the RESV for the primary path) along with the constraint to exclude the protected ABR node 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 node signals the dynamic bypass and will automatically include the XRO with the address of the protected ABR node and propagate 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 object in the Path message.
- If a node-protect dynamic bypass cannot be found or signaled, the PLR node attempts a link-protect dynamic bypass LSP. As in existing implementation of dynamic bypass within the same area, the PLR attempts in the background to signal a node-protect bypass at the receipt of every third Resv refresh message for the primary path.
- Refresh reduction over dynamic bypass will only work if the node-id RRO also contains the interface address. Otherwise the neighbor will not be created once the bypass is activated by the PLR node. The Path state will then time out after three refreshes following the activation of the bypass backup LSP.

Note that a one-to-one detour backup LSP cannot be used at the PLR for the protection of the ABR node. As a result, a PLR node will not signal a one-to-one detour LSP for ABR protection. In addition, an ABR node 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 regardless if the **cspf-on-loose** option is enabled or not on the node. In other words, the router as a transit ABR for the detour path will reject the signaling of an inter-area detour backup LSP.

2.6.13 Automatic Creation of a RSVP Mesh LSP

2.6.13.1 Feature Configuration

The user first creates an LSP template of type mesh P2P:

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

Inside the template the user configures the common LSP and path level parameters or options shared by all LSPs using this template.

Then the user references the peer prefix list which is defined inside a policy statement defined in the global policy manager.

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

The user can associate multiple templates with same or different peer prefix policies. Each application of an LSP template with a given prefix in the prefix list will result 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. This feature does not support the automatic signaling of a secondary path for an LSP. If the user requires the signaling of multiple LSPs to the same destination node, he/she must apply a separate LSP template to the same or different prefix list which contains the same destination node. Each instantiated LSP will have a unique LSP-id and a unique tunnel-ID. This feature also does not support the signaling of a non-CSPF LSP. The selection of the '**no cspf**' option in the LSP template is thus blocked.

Up to 5 peer prefix policies can be associated with a given LSP template at all times. Each time the user executes the above command, with the same or different prefix policy associations, or the user changes a prefix policy associated with an LSP template, the system re-evaluates the prefix policy. The outcome of the re-evaluation will tell MPLS if an existing LSP needs to be torn down or a new LSP needs to be signaled to a destination address which is already in the TE database.

If a /32 prefix is added to (removed from) or if a prefix range is expanded (shrunk) in a prefix list associated with a LSP template, the same prefix policy re-evaluation described above is performed.

The user must perform a **no shutdown** of the template before it takes effect. Once a template is in use, the user must shutdown the template before effecting any changes to the parameters except for those LSP parameters for which the change can be handled with the Make-Before-Break (MBB) procedures. These parameters are **bandwidth** and enabling **fast-reroute** with or without the **hop-limit** or **node-protect** options. For all other parameters, the user shuts down the template and once it is added, removed or modified, the existing instances of the LSP using this template are torn down and re-signaled.

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

2.6.13.2 Feature Behavior

Whether the prefix list contains one or more specific /32 addresses or a range of addresses, an external trigger is required to indicate to MPLS to instantiate an LSP to a node which address matches an entry in the prefix list. The objective of the feature is to provide an automatic creation of a mesh of RSVP LSP to achieve automatic tunneling of LDP-over-RSVP. The external trigger is when the router with the router-id matching an address in the prefix list appears in the Traffic Engineering database. In the latter case, the TE database provides the trigger to MPLS which means this feature operates with CSPF LSP only.

Each instantiation of an LSP template results in RSVP signaling and installing state of a primary path for the LSP to the destination router. The auto- LSP is installed in the Tunnel Table Manager (TTM) and is available to applications such as LDP-over-RSVP, resolution of BGP label routes, resolution of BGP, IGP, and static routes. The auto-LSP can also be used for auto-binding by a VPRN service. The auto-LSP is however not available to be used in a provisioned SDP for explicit binding by services. A consequence of this is that an auto-LSP can also not be used directly for auto-binding of a PW template with the **use-provisioned-sdp** option in BGP-AD VPLS, or FEC129 VLL service. However, an auto-binding of a PW template to an LDP LSP, which is then tunneled over an RSVP auto-LSP is supported.

If the user changes the **bandwidth** parameter in the LSP template, an MBB is performed for all LSPs using the template. If however the **auto-bandwidth** option was enabled in the template, the bandwidth **parameter** change will be saved but will only take effect at the next time the LSP bounces or is re-signaled.

Except for the MBB limitations to the configuration parameter change in the LSP template, MBB procedures for manual and timer based re-signaling of the LSP, for TE Graceful Shutdown and for soft pre-emption are supported.

Note that the use of the '**tools perform router mpls update-path**' command with a mesh LSP is not supported.

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

If while the LSP is UP, with the bypass backup path activated or not, the TE database loses the router-id, it will perform an update to MPLS module which will state router-id is no longer in TE database. This will cause MPLS to tear down all mesh LSPs to this router-id. Note however that if the destination router is not a neighbor of the ingress LER and the user shuts down the IGP instance in the destination router, the router-id corresponding to the IGP instance will only be deleted from the TE database in the ingress LER after the LSA/LSP ages out. If the user brought back up the IGP instance before the LSA/LSP aged out, the ingress LER will delete and re-install the same router-id at the receipt of the updated LSA/LSP. In other words, the RSVP LSPs destined to this router-id will get 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.

There is no overall chassis mode restrictions enforced with the mesh LSP feature. If the chassis-mode, network chassis-mode or IOM type requirements for a feature are not met, the configuration of the corresponding command will not be allowed into the LSP template on the system.

2.6.13.3 Multi-Area and Multi-Instance Support

A router which does not have TE links within a given IGP area/level will not have its router-id discovered in the TE database by other routers in this area/level. In other words, an auto-LSP of type P2P mesh cannot be signaled to a router which does not participate in the area/level of the ingress LER.

A mesh LSP can however be signaled using TE links all belonging 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 ISIS or OSPF are configured on a router, each with its own router-id value, the TE database in other routers will be able to discover TE links advertised by each instance. In such a case, an instance of an LSP can be signaled to each router-id with a CSPF path computed using TE links within each instance.

Finally, if multiple instances of ISIS or OSPF are configured on a destination router each with the same router-id value, a single instance of LSP will be signaled from other routers. If the user shuts down one IGP instance, this will be **no op** 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.

2.6.13.4 Mesh LSP Name Encoding and Statistics

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

Session Name: <lsp-name::path-name>, where lsp-name component is encoded as follows:

TemplateName-DestIpv4Address-TunnelId

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

At ingress LER, the user can enable egress statistics for the auto-created mesh LSP by adding the following configuration to the LSP template:

```
config
  router
    [no] mpls
    lsp-template template-name mesh-p2p]
    no lsp-template template-name
    [no] egress-statistics
    accounting-policy policy-id
    no accounting-policy
    no] collect-stats
```

If there are no stat indices available when an LSP is instantiated, the assignment is failed and the egress-statistics field in the show command for the LSP path will be in the operational DOWN state but in admin UP state.

An auto-created mesh LSP can also have ingress statistics enabled on the egress LER as long as the user specifies the full LSP name following the above syntax.

config>router>mpls>ingress-statistics>lsp *lsp-name sender ip-address*

2.6.14 Timer-based Reversion for RSVP-TE LSPs

The following secondary to primary path reversion is supported for RSVP-TE LSPs:

- Configurable timer-based reversion for primary LSP path
- Manual reversion from secondary to primary path

Normally, an RSVP-TE LSP automatically switches back from using a secondary path to the primary path as soon as the primary path recovers. In some deployments, it is useful to delay reversion or allow manual reversion, rather than allowing an LSP to revert to the primary path as soon as it is available. This feature provides a method to manage fail-overs in the network.

If manual reversion is used, a fall-back timer-based mechanism is required in case a human operator fails to execute the switch back to the primary path. This function is also useful to stagger reversion for large numbers of LSPs.

A reversion timer for an LSP is configured using the CLI as follows:

```
config
  router
    [no] mpls
      lsp
        [no] revert-timer <timer-value>
```

When configured, the revert timer is started as soon as a primary path recovers. The LSP does not revert from the currently used secondary path to the primary path until the timer expires. When configured, the revert-timer is used instead of the existing hold timer.

The timer value can be configured in one minute increments, up to 4320 minutes (72 hours). Once a timer has started, it can be modified using this command. If a new value is entered, then the current timer is canceled (without reverting the LSP) and then restarted using the new value.

The **no** form of the command cancels any currently outstanding revert timer and causes the LSP to revert to the primary path if it is up.

If the LSP secondary path fails while the revert timer is still running, the system cancels the revert-timer and the LSP will thus revert to the primary path immediately. A user can manually force an LSP to revert to the primary path while the revert-timer is still running, using the following tools command:

tools>perform>router>mpls revert lsp *lsp-name*

This command forces the early expiry of the revert timer for the LSP. The primary path must be up in order for this command to work.

2.6.15 Automatic Creation of an RSVP One-Hop LSP

2.6.15.1 Feature Configuration

The user first creates an LSP template of type one-hop:

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

Then the user enables the automatic signaling of one-hop LSP to all direct neighbors using the following command:

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

The LSP and path parameters and options supported in a LSP template of type **one-hop-p2p** are that same as in the LSP template of type **mesh-p2p** except for the parameter **from** which is not allowed in a template of type **one-hop-p2p**. The show command for the auto-LSP displays the actual outgoing interface address in the 'from' field.

Finally the auto-created one-hop LSP can be signaled over both numbered and unnumbered RSVP interfaces.

2.6.15.2 Feature Behavior

Although the provisioning model and CLI syntax differ from that of a mesh LSP only by the absence of a prefix list, the actual behavior is quite different. When the above command is executed, the TE database will keep track of each TE link which comes up to a directly connected IGP neighbor which 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. Thus the **auto-lsp** command with the **one-hop** option will result in one or more LSPs signaled to the IGP neighbor.

Only the router-id of the first IGP instance of the neighbor which advertises a TE link will cause the LSP to be signaled. If subsequently 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 associated now with the other router-id.

The state of a one-hop LSP once signaled follows the following behavior:

- If the interface used by the TE link goes down or BFD times out and the RSVP interface registered with BFD, the LSP path moves to the bypass backup LSP if the primary path is associated with one.
- If while the one-hop LSP is UP, with the bypass backup path activated or not, the association of the TE-link with a router-id is removed in the TE databases, the one-hop LSP is torn down. This would be the case if the interface used by the TE link is deleted or if the interface is shutdown in the context of RSVP.
- If while the LSP is UP, with the bypass backup path activated or not, the TE database loses the router-id, it will perform two separate updates to MPLS module. The first one updates the loss of the TE link association which will cause action (B) above for the one-hop LSP. The other update will state router-id is no longer in 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 which 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 is not subject to age-out as for a non-directly connected destination router.

All other feature behavior, limitations, and statistics support are the same as for an auto-LSP of type **mesh-p2p**.

2.6.16 MPLS Entropy Label

The router supports the MPLS entropy label (RFC 6790) on RSVP-TE LSPs used for IGP and BGP shortcuts. This allows LSR nodes in a network to load-balance labeled packets in a much more granular fashion than allowed by simply hashing on the standard label stack. See [MPLS Entropy Label and Hash Label](#) for further information.


2.7 Point-to-Multipoint (P2MP) RSVP LSP

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

2.7.1 Application in Video Broadcast

Figure 36 illustrates the use of the 7750 SR product family in triple play application (TPSDA). The Broadband Service Router (BSR) is a 7750 SR and the Broadband Service Aggregator (BSA) is the 7450 ESS.

Figure 36 Application of P2MP LSP in Video Broadcast



A PIM-free core network can be achieved by deploying P2MP LSPs using other core routers. The router can act as the ingress LER receiving the multicast packets from the multicast source and forwarding them over the P2MP LSP.

A router can act as a leaf for the P2MP LSP tree initiated from the head-end router co-located with the video source. The router can also act as a branch node serving other leaf nodes and supports the replication of multicast packets over P2MP LSPs.

2.7.2 P2MP LSP Data Plane

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

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

When a router sends a packet over a P2MP LSP which egresses on an Ethernet-based network interface, the Ethernet frame uses a MAC unicast destination address when sending the packet over the primary P2MP LSP instance or over a P2P bypass LSP). Note that a MAC multicast destination address is also allowed in the draft-ietf-mpls-multicast-encaps. Thus, at the ingress network interface on an Ethernet port, the router can accept both types of Ethernet destination addresses.

2.7.2.1 Procedures at Ingress LER Node

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

1. First, the P2MP LSP state is established via the control plane. Each leaf of the P2MP LSP will have a next-hop label forwarding entry (NHLFE) configured in the forwarding plane for each outgoing interface.
2. The user maps a specific multicast destination group address to the P2MP LSP in the base router instance by configuring a static multicast group under a tunnel interface representing the P2MP LSP.
3. An FTN entry is programmed at the ingress of the head-end node that maps the FEC of a received user IP multicast packet to a list of outgoing interfaces (OIF) and corresponding NHLFEs.
4. The head-end node replicates the received IP multicast packet to each NHLFE. Replication is performed at ingress toward the fabric and/or at egress forwarding engine depending on the location of the OIF.
5. At ingress, the head-end node performs a PUSH operation on each of the replicated packets.

2.7.2.2 Procedures at LSR Node

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

- The LSR performs a label swapping operation on a leaf of the P2MP LSP. This is a conventional operation of an LSR in a P2P LSP. An ILM entry is programmed at the ingress of the LSR to map an incoming label to a NHLFE.

The following is an exception handling procedure for control packets received on an ILM in an LSR.

- Packets that arrive with the TTL in the outer label expiring are sent to the CPM for further processing and are not forwarded to the egress NHLFE.

2.7.2.3 Procedures at Branch LSR Node

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

- The LSR performs a replication and a label swapping for each leaf of the P2MP LSP. An ILM entry is programmed at the ingress of the LSR to map an incoming label to a list of OIF and corresponding NHLFEs.
- There is a limit of 127 OIF/NHLFEs per ILM entry.

The following is an exception handling procedure for control packets received on an ILM in a branch LSR:

- Packets that arrive with the TTL in the outer label expiring are sent to the CPM for further processing and not copied to the LSP branches.

2.7.2.4 Procedures at Egress LER Node

The following procedures occur at the leaf node of the P2MP LSP (egress LER):

- The egress LER performs a pop operation. An ILM entry is programmed at the ingress of the egress LER to map an incoming label to a list of next-hop/OIF.

The following is an exception handling procedure for control packets received on an ILM in an egress LER.

- The packet is sent to the CPM for further processing if there is any of the IP header exception handling conditions set after the label is popped: 127/8 destination address, router alert option set, or any other options set.

2.7.2.5 Procedures at BUD LSR Node

The following are procedures at an LSR node which is both a branch node and an egress leaf node (bud node):

- The bud LSR performs a pop operation on one or many replications of the received packet and a swap operation of the remaining replications. An ILM entry is programmed at ingress of the LSR to map the incoming label to list of NHLFE/OIF and next-hop/OIF.

Note however, the exact same packets are replicated to an LSP leaf and to a local interface.

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

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

2.7.3 Ingress Path Management for P2MP LSP Packets

The SR OS provides the ingress multicast path management (IMPM) capability that allows users to manage the way IP multicast streams are forwarded over the router's fabric and to maximize the use of the fabric multicast path capacity.

IMPM consists of two components, a bandwidth policy and a multicast information policy. The bandwidth policy configures the parameters of the multicast paths to the fabric. This includes the rate limit and the multicast queue parameters of each path. The multicast information policy configures the bandwidth and preference parameters of individual multicast flows corresponding to a channel, for example, a <*,G> or a <S,G>, or a bundle of channels.

By default, the XCM (on the 7950 XRS) and the IOM-2 and IOM-3/IMM (on the 7750 SR and 7450 ESS) ingress data paths provides two multicast paths through the fabric referred to as high-priority path and low-priority path respectively. When a multicast packet is received on an ingress network or access interface or on a VPLS SAP, the packet's classification will determine its forwarding class and priority or profile as per the ingress QoS policy. This then determines which of the SAP or interface multicast queues it must be stored in. By default SAP and interface expedited forwarding class queues forward over the high-priority multicast path and the non expedited forwarding class queues forward over the low-priority multicast path.

When IMPM on the ingress MDA is enabled on the 7750 SR and 7450 ESS, one or more multicast paths are enabled depending on the IOM type. In addition, for all routers, multicast flows managed by IMPM will be stored in a separate shared multicast queue for each multicast path. These queues are configured in the bandwidth policy.

IMPM maps a packet to one of the paths dynamically based on monitoring the bandwidth usage of each packet flow matching a $\langle *, G \rangle$ or $\langle S, G \rangle$ record. The multicast bandwidth manager also assigns multicast flows to a primary path, and ancillary path for IOM-2, based on the flow preference until the rate limits of each path is reached (this does not apply to the 7950 XRS). At that point in time, a multicast flow is mapped to the secondary flow. If a path congests, the bandwidth manager will remove and black-hole lower preference flows to guarantee bandwidth to higher preference flows. The preference of a multicast flow is configured in the multicast info policy.

A packet received on a P2MP LSP ILM is managed by IMPM when IMPM is enabled on the ingress XMA or the ingress MDA and the packet matches a specific multicast record. When IMPM is enabled but the packet does not match a multicast record, or when IMPM is disabled, a packet received on a P2MP LSP ILM is mapped to a multicast path; the mapping is different on the 7750 SR and 7450 ESS and depends on whether the ingress IOM is an IOM-2 or IOM-3.

2.7.3.1 Ingress P2MP Path Management on XCM/IOM-3/IMMs

On an ingress XCM or IOM-3/IMM, there are multiple multicast paths available to forward multicast packets, depending on the hardware being used. Each path has a set of multicast queues and associated with it. Two paths are enabled by default, a primary path and a secondary path, and represent the high-priority and low-priority paths respectively. Each VPLS SAP, access interface, and network interface will have a set of per forwarding class multicast and/or broadcast queues which are defined in the ingress QoS policy associated with them. The expedited queues will be attached to the primary path while the non-expedited queues will be attached to secondary path.

When IMPM is enabled and/or when a P2MP LSP ILM exists on the ingress XCM or IOM-3/IMM, the remaining multicast paths are also enabled. 16 multicast paths are supported by default with 28 on 7950 XRS systems and 7750 SR-12e systems, with the latter having the **tools perform system set-fabric-speed fabric-speed-b**. One path remains as a secondary path and the rest are primary paths.

A separate pair of shared multicast queues is created on each of the primary paths, one for IMPM managed packets and one for P2MP LSP packets not managed by IMPM. The secondary path does not forward IMPM managed packets or P2MP LSP packets. These queues have default rate (PIR=CIR) and CBS/MBS/Hi-Priority-Only thresholds but can be changed away from default under the bandwidth policy.

A VPLS snooped packet, a PIM routed packet, or a P2MP LSP packet is managed by IMPM if it matches a <*,G> or a <S,G> multicast record in the ingress forwarding table and IMPM is enabled on the ingress XMA or the MDA where the packet is received. The user enables IMPM on the ingress XMA data path or the MDA data path using the **config>card>mda>ingress>mcast-path-management** command.

A packet received on an IP interface and to be forwarded to a P2MP LSP NHLFE or a packet received on a P2MP LSP ILM is not managed by IMPM when IMPM is disabled on the ingress XMA or the MDA where the packet is received or when IMPM is enabled but the packet does not match any multicast record. A P2MP LSP packet duplicated at a branch LSR node is an example of a packet not managed by IMPM even when IMPM is enabled on the ingress XMA or the MDA where the P2MP LSP ILM exists. A packet forwarded over a P2MP LSP at an ingress LER and which matches a <*,G> or a <S,G> is an example of a packet which is not managed by IMPM if IMPM is disabled on the ingress XMA or the MDA where the packet is received.

When a P2MP LSP packet is not managed by IMPM, it is stored in the unmanaged P2MP shared queue of one of the primary multicast paths.

By default, non-managed P2MP LSP traffic is distributed across the IMPM primary paths using hash mechanisms. This can be optimized by enabling IMPM on any forwarding complex, which allows the system to redistributed this traffic on all forwarding complexes across the IMPM paths to achieve a more even capacity distribution. Be aware that enabling IMPM will cause routed and VPLS (IGMP and PIM) snooped IP multicast groups to be managed by IMPM.

The above ingress data path procedures apply to packets of a P2MP LSP at ingress LER, LSR, branch LSR, bud LSR, and egress LER. Note that in the presence of both IMPM managed traffic and unmanaged P2MP LSP traffic on the same ingress forwarding plane, the user must account for the presence of the unmanaged traffic on the same path when setting the rate limit for an IMPM path in the bandwidth policy.

2.7.3.2 Ingress P2MP Path Management on IOM-2

The following procedures apply at the ingress data path for packets received from or to be forwarded to a P2MP LSP at ingress LER, LSR, branch LSR, bud LSR, and egress LER.

On ingress IOM-2, there are 3 multicast paths which are available for forwarding multicast packets. Each path has a set of multicast queues and a multicast VoQ associated with it. Two paths are enabled by default, a primary path and a secondary path, and represent the high-priority and low-priority paths respectively. Each VPLS SAP, access interface, and network interface will have a set of per forwarding class multicast and/or broadcast queues which are defined in the ingress QoS policy associated with them. The expedited queues will be attached to the primary path while the non-expedited queues will be attached to the secondary path.

When IMPM is disabled, packets of P2MP LSP arriving on a network interface will be queued in that interface queue corresponding to the forwarding class of the packet.

When the user enables IMPM on the ingress MDA, a third multicast path, referred to as ancillary path, is added on the ingress IOM-2. This path reuses unused capacity from the unicast paths. The high-priority and low-priority paths are renamed as primary and secondary paths respectively.

A VPLS snooped packet or a PIM routed packet is managed by IMPM if it matches a $\langle *, G \rangle$ or a $\langle S, G \rangle$ multicast record in the ingress IOM-2 forwarding table and IMPM is enabled on the ingress MDA where the packet is received. The user enables IMPM on the ingress MDA data path using the **config>card>mda>ingress>mcast-path-management** command.

A P2MP LSP packet which matches a multicast record is also managed by IMPM on ingress IOM-2 and is thus distributed to one of the primary, ancillary, or secondary path according to the congestion level of the paths and the preference of the packet's multicast flow as configured in the multicast info policy 2.

2.7.4 RSVP Control Plane in a P2MP LSP

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

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

A P2MP LSP is identified by the combination of $\langle \text{P2MP ID, tunnel ID, extended tunnel ID} \rangle$ part of the P2MP session object, and $\langle \text{tunnel sender address, LSP ID} \rangle$ fields in the P2MP sender_template object.

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

The following are characteristics of this feature:

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

- The **to** parameter is not available at the LSP level but at the path level of each S2L sub-LSP of the primary or secondary instance of this P2MP LSP.
- The hold-timer configured in the **config>router>mpls>hold-timer** context applies when signaling or re-signaling an individual S2L sub-LSP path. It does not apply when the entire tree is signaled or re-signaled.
- The head-end node can add and/or remove a S2L sub-LSP of a specific leaf node without impacting forwarding over the already established S2L sub-LSPs of this P2MP LSP and without re-signaling them.
- The head-end node performs a make-before break (MBB) on an individual S2L path of a primary P2MP instance whenever it applies the FRR global revertive procedures to this path. If CSPF finds a new path, RSVP signals this S2L path with the same LSP-ID as the existing path.
- All other configuration changes, such as adaptive/no-adaptive (when an MBB is in progress), use-te-metric, no-frr, cspf/no-cspf, result in the tear-down and re-try of all affected S2L paths.
- MPLS requests CSPF to re-compute the whole set of S2L paths of a given active P2MP instance each time the P2MP re-signal timer expires. The P2MP re-signal timer is configured separately from the P2P LSP. MPLS performs a global MBB and moves each S2L sub-LSP in the instance into its new path using a new P2MP LSP ID if the global MBB is successful. This is regardless of the cost of the new S2L path.
- MPLS will request CSPF to re-compute the whole set of S2L paths of a given active P2MP instance each time the user performs a manual re-signal of the P2MP instance. MPLS then always performs a global MBB and moves each S2L sub-LSP in the instance into its new path using a new P2MP LSP ID if the global MBB is successful. This is regardless of the cost of the new S2L path. The user executes a manual re-signal of the P2MP LSP instance using the command:
tools>perform>router>mpls>resignal p2mp-lsp *lsp-name* p2mp-instance *instance-name*.
- When performing global MBB, MPLS runs a separate MBB on each S2L in the P2MP LSP instance. If an S2L MBB does not succeed the first time, MPLS will re-try the S2L using the re-try timer and re-try count values inherited from P2MP LSP configuration. However, there will be a global MBB timer set to 600 seconds and which is not configurable. If the global MBB succeeds, for example, all S2L MBBs have succeeded, before the global timer expires, MPLS moves the all S2L sub-LSPs into their new path. Otherwise when this timer expires, MPLS checks if all S2L paths have at least tried once. If so, it then aborts the global MBB. If not, it will continue until all S2Ls have re-tried once and then aborts the global MBB. Once global MBB is aborted, MPLS will move all S2L sub-LSPs into the new paths only if the set of S2Ls with a new path found is a superset of the S2Ls which have a current path which is up.

- While make-before break is being performed on individual S2L sub-LSP paths, the P2MP LSP will continue forwarding packets on S2L sub-LSP paths which are not being re-optimized and on the older S2L sub-LSP paths for which make-before-break operation was not successful. MBB will thus result in duplication of packets until the old path is torn down.
- The MPLS data path of an LSR node, branch LSR node, and bud LSR node will be able to re-merge S2L sub-LSP paths of the same P2MP LSP in case their ILM is on different incoming interfaces and their NHLFE is on the same or different outgoing interfaces. This could occur anytime there are equal cost paths through this node for the S2L sub-LSPs of this P2MP LSP.
- Link-protect FRR bypass using P2P LSPs is supported. In link protect, the PLR protecting an interface to a branch LSR will only make use of a single P2P bypass LSP to protect all S2L sub-LSPs traversing the protected interface.
- Refresh reduction on RSVP interface and on P2P bypass LSP protecting one or more S2L sub-LSPs.
- A manual bypass LSP cannot be used for protecting S2L paths of a P2MP LSP.
- The following MPLS features do operate with P2MP LSP:
 - BFD on RSVP interface.
 - MD5 on RSVP interface.
 - IGP metric and TE metric for computing the path of the P2MP LSP with CSPF.
 - SRLG constraint for computing the path of the P2MP LSP with CSPF. SRLG is supported on FRR backup path only.
 - TE graceful shutdown.
 - Admin group constraint.
- The following MPLS features are not operable with P2MP LSP:
 - Class based forwarding over P2MP RSVP LSP.
 - LDP-over-RSVP where the RSVP LSP is a P2MP LSP.
 - Diff-Serv TE.
 - Soft preemption of RSVP P2MP LSP.

2.7.5 Forwarding Multicast Packets over RSVP P2MP LSP in the Base Router

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

2.7.5.1 Procedures at Ingress LER Node

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

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

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

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

A multicast packet which is received on an interface and which succeeds the RPF check for the source address will be replicated and forwarded to all OIFs which correspond to the branches of the P2MP LSP. The packet is sent on each OIF with the label stack indicated in the NHLFE of this OIF. The packets will also be replicated and forwarded natively on all OIFs which have received IGMP or PIM joins for this <S,G>.

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

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

2.7.5.2 Procedures at Egress LER Node

2.7.5.2.1 Procedures with a Primary Tunnel Interface

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

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

The egress LER accepts multicast packets using the following methods:

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

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

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

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

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

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

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

2.8 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 segment 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 Segment ID (SID).

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

Segment routing using MPLS labels can be used in both shortest path routing applications (refer to the Routing Protocols guide for more 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- or source-routed forwarding path with the same mechanism; re-use IGP to minimize the number of control plane protocols

- 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
- FRR: expand coverage of basic LFA to any topology with the use of source-routed backup path; pre-computation and set up of backup path without additional signaling
- support LFA policies with shared-risk constraints, admin-groups, link/node protection
- Traffic Engineering (TE) should include loose/strict options, distributed and centralized TE, path disjointness, ECMP-awareness, limited or no per-service state on midpoint and tail-end routers
- 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)

2.8.1 SR-TE Support

The following MPLS commands and nodes are supported:

- Global MPLS-level commands and nodes:
interface, lsp, path, shutdown
- LSP-level commands and nodes:
bgp-shortcut, bgp-transport-tunnel, exclude, hop-limit, include, metric, primary, retry-limit, retry-timer, shutdown, to, from, vprn-auto-bind
- Only a primary path is supported with a SR-TE LSP. The following primary path level commands and nodes are supported with SR-TE LSP:
bandwidth, exclude, hop-limit, include, priority, shutdown

The following MPLS commands and nodes are *not* supported:

- Global MPLS level commands and nodes not applicable to SR-TE LSP (configuration is ignored):

admin-group-frr, auto-bandwidth-multipliers, auto-lsp, bypass-resignal-timer, cspf-on-loose-hop, dynamic-bypass, exponent-backoff-retry, frr-object, hold-timer, ingress-statistics, least-fill-min-thd, least-fill-reoptim-thd, logger-event-bundling, lsp-init-retry-timeout, lsp-template, max-bypass-associations, mbb-prefer-current-hops, mpls-tp, p2mp-resignal-timer, p2mp-s2l-fast-retry, p2p-active-path-fast-retry, retry-on-igp-overload, secondary-fast-retry-timer, shortcut-local-ttl-propagate, shortcut-transit-ttl-propagate, srlg-database, srlg-frr, static-lsp, static-lsp-fast-retry, user-srlg-db

- LSP level commands and nodes not supported with SR-TE LSP (configuration blocked):

adaptive, adspec, auto-bandwidth, bfd-enable, class-forwarding, class-type, dest-global-id, dest-tunnel-number, egress-statistics, exclude-node, fast-reroute, ingress-statistics, ldp-over-rsvp, least-fill, load-balancing-weight, main-ct-retry-limit, p2mp-id, primary-p2mp-instance, propagate-admin-group, protect-tp-path, revert-timer, protect-tp-path, rsvp-resv-style, secondary, working-tp-path

- LSP-level commands and nodes not supported with SR-TE LSP (configuration ignored):

igp-shortcut, cspf

- The following primary path level commands and nodes are not supported with SR-TE LSP (configuration blocked):

adaptive, backup-class-type, bfd-enable, class-type, record, record-label

- Secondary path is not supported.

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** and **primary** 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](#).

2.8.2 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 Traffic Engineering controller referred to as a Path Computation Element (PCE). This feature works with the Nokia stateful PCE which is part of the Network Services Platform (NSP).

The SR OS support three different modes of operations configurable on a per SR-TE LSP basis:

- When the path of the LSP is computed by the router acting as a PCE Client (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
- Does not support CSPF and local path computation takes the form of hop-to-label translation
- The capability exists to report a 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. In other words, 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 **pce-computation** option for the SR-TE LSP so PCE can perform path computation at the request of the PCC only. PCC retains control.

The capability exists to report a 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 PCE can perform path computation and updates following a network event without the explicit request from the PCC. 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 PCE for a specific LSP using the **pce-computation** and **pce-control** commands.

The **pce-computation** 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 pce-computation**. Enabling **pce-computation** requires that the **cspf** option is also enabled, otherwise the command is rejected. If the **cspf** option is disabled for an LSP, the **pce-computation** option will also be automatically disabled.

Enabling **cspf** without enabling **pce-computation** for a SR-TE LSP means that, internally, the router still performs label translation as if **cspf** was disabled, because there is no support of CSPF for a SR-TE LSP on the router.

The **pce-control** option allows the router to delegate full control of the LSP to the PCE (PCE-controlled). Enabling it means the PCE is acting in Active Stateful mode for this LSP and allows PCE to reroute the path following a failure or to re-optimize the path and update the router without requiring the router to request it.

**Note:**

- The user can delegate CSPF and non-CSPF LSPs.
- The user can delegate LSPs which have the **pce-computation** option enabled or disabled. The LSP maintains its latest active path computed by 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 re-optimization. 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 given type. The PCC reports both CSPF and non-CSPF LSP. The default value is disabled (**no pce-report**). This default value controls the introduction of 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 reporting LSP to PCE. The default value is to inherit the global MPLS level value. The **inherit** value returns 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.

2.8.2.1 PCC-Initiated and PCC-Controlled LSP

In this mode of operation, the user configures the LSP name and the primary 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.

To configure the primary 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 an address of a loopback address, it will be translated into an adjacency SID as explained below and thus does not guarantee that the same specific TE link is picked.

To use an SR-TE path that consists of unprotected adjacency SIDs, then 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 in an RSVP-TE LSP. These IDs will be useful to represent the MBB path of the same SR-TE LSP which need to co-exist during the update of the primary path.



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

The router retains full control of the path of the LSP. CSPF is not supported and, as such, 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-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:

Step 1. MPLS passes the path information to TE-DB, which converts the list of hops into a label stack by scanning the TE database for adjacency and node SID information which belongs to the router or link identified by each hop address. If the conversion is successful, the TE database will return the actual selected hop SIDs plus labels as well the configured path hop addresses which 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 database, 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, then 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 which 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 in the case of a LAN interface. In this case, 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 OSPF instances; not only the IGP instance which resolved the prefix of the destination address of the LSP in RTM is used. For the first instance via which 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 thus taken from the same IGP area and instance.
- Unnumbered network IP interfaces, which are supported in the router's TE database, can be selected when converting the hops into an adjacency SID label when the user has entered the address of a loopback interface as a strict hop; however, the user cannot configure an unnumbered interface as a hop in the path definition.



Note: For the hop-to-label translation to operate, the user must enable TE on the network links, meaning to add the network interfaces to MPLS and RSVP. In addition, the user must enable the **traffic-engineering** option on all participating router IGP instances. Note that if any router has the **database-export** option enabled in the participating IGP instances to populate the learned IGP link state information into the TE-DB, then enabling of the **traffic-engineering** option is not required. For consistency purposes, it is recommended to have the **traffic-engineering** option always enabled.

- Step 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 some downstream router, then the node SID label is pushed.

In both cases, the SR-TE LSP tracks and rides the SR shortest path tunnel of the SID of the first hop.

- Step 3.** In the case where the router is configured as a PCC and has a PCEP session to a PCE, the router sends a PCRpt message to update PCE with the state of UP and the RRO object for each LSP which has the **pce-report** option enabled. 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.

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

2.8.2.1.1 Guidelines for Using PCC-Initiated and PCC-Controlled LSPs


In R14, the router does not support CSPF path computation for a SR-TE LSP and uses the hop-to-label translation to compute the path. Thus the ingress LER does not monitor network events which affect the reachability of the adjacency SID or node SID used in the label stack of the LSP, and thus the label stack is not updated to reflect changes in the path. As a result, it is 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

In addition, the user can configure a SR-TE LSP with a single loose-hop using the anycast SID concept to provide LSR node protection within a given plane of the network TE topology. This is illustrated in Figure 37. The user configures all LSRs in a given plane with the same loopback interface address, which must be different from that of the system interface and the router-id of the router, and assigns them the same node-SID index value. All routers must use the same SRGB.

Figure 37 Multi-plane Traffic Engineering with Node Protection



Then user configures in a LER a SR-TE LSP to some destination and adds to its path a loose-hop matching the anycast loopback address. The SR-TE LSP to any destination will hop over the closest of the LSRs owning the anycast SID because the resolution of the node-SID for that anycast loopback address uses the closest router. When that router fails, the resolution is updated to the next closest router owning the anycast SID without changing the label stack of the SR-TE LSP.

2.8.2.2 PCC-Initiated and PCE-Computed/Controlled LSP

In this mode of operation, the ingress LER uses Path Computation Element Communication Protocol (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 PCE Client (PCC).

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

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

The following procedure describes configuring and programming 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 via CLI or via OSS/SAM.
The configuration dictates which PCE control mode is desired: active (**pce-control** option enabled) or passive (**pce-computation** 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 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 PCE, the PCE ignores them and computes the path with the other constraints, excepting 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. 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. 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: 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 IGP. This causes IGP to flood the link SRLG and admin-group membership in its participating area and for the PCE to learn it in its TE database.

6. The PE router updates the CPM and the 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. PE router sends a PCRpt message to update PCE with the state of UP 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 re-optimization, 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 CPM and the data path with the new path.

10. The PE router sends a new PCRpt message to update PCE with the state of UP 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 PCC to 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 a new path programmed into the data path, the LSP databases of PCC and PCE to be synchronized, and the delegation to be returned to PCE.

In the case of an SR-TE LSP, MBB is not supported. Therefore, PCC first tears down the LSP and sends a PCRpt message to 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 followed for an LSP with an active path:

- If the user enabled the **pce-computation** option on a PCC-controlled LSP which has 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 procedures above are followed.
- If the user enabled the **pce-control** option on a PCC-controlled or PCE-computed LSP which has an active path, PCC will issue a PCRpt message to PCE with the state of UP and the RRO of the active path. It will set delegation control flag to delegate control to PCE. PCE will keep the active path of the LSP and will not update until the next network event or re-optimization. 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 PCE, the user must configure the PCE path profile ID and path group ID the LSP belongs to. These parameters are passed transparently by PCC to PCE and are thus 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 PCE 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 section of this guide.

2.8.3 SR-TE LSP Path Computation

For PCC-controlled SR-TE LSPs, CSPF is not supported on the router. Whether the **cspf** option is enabled or disabled for a SR-TE LSP, 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 LSP](#) 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 **pce-computation** option, as explained in [SR-TE LSP Instantiation](#). The user can further delegate the re-optimization of the LSP to the PCE by enabling the **pce-control** option. In both cases, 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 which the ingress LER can push for a given 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 so as to allow room to insert additional transport, service, and other labels when packets are forwarded in a given 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 1-11 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 0-3 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](#) and always signaled by 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 PCC to PCE in the Segment-ID (SID) Depth value in a METRIC object for both a PCE-computed LSP and a PCE-controlled LSP. 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 PCC.

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

2.8.3.1 Service and Shortcut Application SR-TE Label Stack Check

In Release 14.0.R1, if a packet forwarded in a service or a shortcut application resulted in the net label stack size being pushed on the packet to exceed the maximum label stack supported by the router, the packet got dropped on the egress. In Release 14.0.R4, 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.

To that effect, 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 will then add the additional, context-specific labels, such as service label, entropy/hash label, and control-word, required to forward the packet in that context and to check 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, it 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 a SDP with one or more SR-TE LSP names, 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, a new spoke-sdp flag will be displayed in the show output of the service: "labelStackLimitExceeded". Similarly, 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 will become unresolved.

The value of **additional-frr-labels** *labels* is checked against the maximum value across all IGP instances of the parameter *frr-overhead*. This parameter is computed within a given IGP instance as shown in [Table 12](#).

Table 12 frr-overhead Parameter Values

Condition	frr-overhead Parameter Value
segment-routing is disabled in the IGP instance	0
segment-routing is enabled but remote-lfa is disabled	0

Table 12 **frr-overhead Parameter Values (Continued)**

Condition	frr-overhead Parameter Value
segment-routing is enabled and remote-lfa is enabled	1

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

Example:

1. The user configures the **config>router>isis>loopfree-alternate remote-lfa** command.
2. The user creates a new SR-TE LSP or changes the configuration of an existing as follows: **mpls>lsp>max-sr-labels 10 additional-frr-labels 0**.
3. 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 then adds **max-sr-labels** and **additional-frr-labels** and checks that it is lower 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 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*, then IGP will check that all SR-TE LSPs can properly account for the overhead or the change is rejected. On the IGP, enabling **remote-lfa** may cause the *frr-overhead* to change.

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 has the loopfree-alternate command disabled.
- The user attempts to configure **isis>loopfree-alternate remote-lfa**. This changes *frr-overhead* to 1. This configuration change would be blocked.

2.8.4 SR-TE LSP Protection

Each path is locally protected along the network using LFA/remote-LFA next-hop whenever possible. The protection of a node SID re-uses the LFA and remote LFA features introduced with segment routing shortest path tunnels; the protection of an adjacency SID has been added to the SR OS in the specific context of an SR-TE LSP to augment the protection level. The user must enable the **loopfree-alternate [remote-lfa]** option 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 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 a 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 of 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:

- When the ILM record is for a node SID of a downstream router which 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 LER and LSR. Thus, 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 LFA.

- When the ILM is for a node SID of a directly connected router, then 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, as such, can only protect the link to the neighboring router.

- When 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 (as above).

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 this node SID is reachable over parallel links. The selection is based on lowest interface ID. When no parallel links exist, then regular LFA/rLFA 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 thus 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 itself will point 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. Thus, the FRR link or node protection at ingress LER is inherited directly from the SR shortest path tunnel.

When an adjacency to a neighbor fails, 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 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 seconds.

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

- The network IP interface went down due 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 deleted an interface in the **config>router>ospf/isis** context.


Note:

- The above 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 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 re-programmed if the adjacency is restored before the above 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 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 which 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-alternate** option in ISIS 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 a SR-TE LSP. In that case, the user can disable protection for all adjacency SIDs formed over a given 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 PCE the constraint to use only adjacency SIDs, which are not protected. The Path Profile ID is configured in PCE with the no-protection constraint.

2.8.5 Static Route Resolution using SR-TE LSP

The user can forward packets of a static route to an indirect next-hop over an SR-TE LSP programmed in TTM by configuring 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 name1]
[no] [lsp name2]
.
.
[no] [lsp name-N]
exit
[no] disallow-igp
```

```
        exit
    exit
```

The user can select the **sr-te** tunnel type and either specify a list of SR-TE LSP names to use to reach the indirect next-hop of this static route or have the SR-TE LSPs automatically select the indirect next-hop in TTM.

2.8.6 BGP Shortcuts Using SR-TE LSP

The user can forward packets of BGP prefixes over an SR-TE LSP programmed in TTM by configuring the following BGP shortcut tunnel binding command:

CLI Syntax:

```
config>router>bgp>next-hop-resolution
    shortcut-tunnel
        [no] family {ipv4}
            resolution {any | disabled | filter}
            resolution-filter
                [no] sr-te
        exit
    exit
exit
```

2.8.7 BGP Label Route Resolution Using SR-TE LSP

The user can enable SR-TE LSP, as programmed in TTM, for resolving the next-hop of a BGP IPv4 or IPv6 (6PE) label route by enabling the following BGP transport tunnel command:

CLI Syntax:

```
config>router>bgp>next-hop-res>
    label-route-transport-tunnel
        [no] family {ipv4, vpn, ipv6}
            resolution {any | disabled | filter}
            resolution-filter
                [no] sr-te
        exit
    exit
exit
```

2.8.8 Service Packet Forwarding using SR-TE LSP

An SDP sub-type of the MPLS encapsulation type allows service binding to a SR-TE LSP programmed in TTM by MPLS:

Example: *A:7950 XRS-20# configure service sdp 100 mpls create
 *A:7950 XRS-20>config>service>sdp\$ sr-te-lsp lsp-name

The user can specify up to 16 SR-TE LSP names. The destination address of all LSPs must match that of the SDP far-end option. Service packets are sprayed over the set of LSPs in the SDP using the same procedures used for tunnel selection in ECMP. Since the SR-TE LSP feature does not support ECMP at the ingress LER when the outer SR tunnel is a node-SID with multiple next-hops, the first next-hop of each of the 16 LSPs is used for the service packet spraying. 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 a SR-TE LSP can be configured to static (**off**), T-LDP (**tldp**), or BGP (**bgp**).

An SR-TE LSP can be used in VPRN auto-bind 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 command.

This auto-bind 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

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

- VLL, LDP VPLS, IES/VPRN spoke-interface, R-VPLS, BGP EVPN

- BGP-AD VPLS, BGP-VPLS, BGP VPWS when the **use-provisioned-sdp** option is enabled in the binding to the PW template
- intra-AS BGP VPRN for VPN-IPv4 and VPN-IPv6 prefixes with both auto-bind and explicit SDP
- inter-AS options B and C for VPN-IPv4 and VPN-IPv6 VPRN prefix resolution
- IPv4 BGP shortcut and IPv4 BGP label route resolution
- IPv4 static route resolution
- multicast over IES/VPRN spoke interface with **spoke-sdp** riding a SR-TE LSP

2.8.9 Data Path Support

The support of SR-TE in the data path requires that the ingress LER pushes 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/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/loose hop.

This actually represents the NHLFE of the SR shortest path tunnel to the first strict/loose hop. SR OS keeps the SR shortest path tunnel to a downstream node SID or adjacency SID in the tunnel table and thus its NHLFE is readily available. The rest of the label stack is not meaningful to the forwarding decision. In this document, “super NHLFE” refers to this part of the label stack because it can have a much larger size.

As a result, an SR-TE LSP is modeled in the ingress LER data path as a hierarchical LSP with the super NHLFE is tunneled over the NHLFE of the SR shortest path tunnel to the first strict/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 riding the same SR shortest path tunnel, and will thus consume each one super NHLFE but they will be pointing to a single NHLFE, or set of NHLFEs when ECMP exists for the first strict/loose hop, of the first hop SR tunnel.

Also, 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 when the user configured an empty path SR-TE LSP which uses the router's hop-to-label translation. In that 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.
- The SR-TE LSP feature does not currently support ECMP at the ingress LER when the outer SR tunnel is a node-SID with multiple next-hops. The data path will program the first next-hop of the node SID for the SR-TE LSP. This limitation does not exist at an LSR.
- If the first hop SR tunnel, node or adjacency SID, goes down the SR module informs MPLS that outer tunnel down and MPLS brings the SR-TE LSP down and requests SR to delete the SR-TE LSP in IOM.

The data path behavior at LSR and egress LER for an SR-TE LSP is similar to that of 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, then 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, service, hash, 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 maximum value of 11 is obtained for an SR-TE LSP whose path is not protected via FRR backup and with no entropy or hash label feature enabled when such an LSP is used as a shortcut for an IGP IPv4/IPv6 prefix or as a shortcut for BGP IPv4/IPv6. In this case, the IPv6 prefix requires pushing the IPv6 explicit-null label at the bottom of the stack. This leaves 11 labels for the SR-TE LSP.

The default value of 6 is obtained in the worst cases, such as forwarding a vprn-ping packet for an inter-AS VPN-IP prefix in Option C:

6 SR-TE labels + 1 remote LFA SR label + BGP 3107 label + ELI (RFC 6790) + EL (entropy label) + service label + OAM Router Alert label = 12 labels.

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

LER role:


- Push up to 12 labels.
- Pop up to 8 labels of which 4 labels can be transport labels

LSR role:

- Pop up to 5 labels and swap one label for a total of 6 labels
- LSR hash of a packet with up to 16 labels

An example of the label stack pushed by the ingress LER and by a LSR acting as a PLR is illustrated in [Figure 38](#).

Figure 38 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:** 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 pre-program 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 as it only deals with the programming of the ILM and primary/backup NHLFE of its adjacency SIDs and its local and remote node SIDs.

VPNRN service in node A forwards a packet to the VPN-IPv4 prefix X advertised by BGP peer D. [Figure 38](#) shows the resulting data path at each node for the primary path and for the FRR backup path at LSR B.

2.8.9.1 SR-TE LSP Metric and MTU Settings

The MPLS module assigns a TE-LSP the maximum LSP metric value of 16777215 when the local router provides the hop-to-label translation for its path. For a TE-LSP that uses PCE for path computation (**pce-computation** option enabled) by PCE and/or which has its control delegated to PCE (**pce-control** enabled), the latter will return the computed LSP IGP or TE metric in the PCReq and PCUpd messages. In both cases, the user can override the returned value by configuring an admin metric using the command **config>router>mpls>lsp>metric**.

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

The following are the details of this calculation:

$$\text{SR\_Tunnel\_MTU} = \text{MIN} \{ \text{Cfg\_SR\_MTU}, \text{IGP\_Tunnel\_MTU} - (1 + \text{frr-overhead}) * 4 \}$$

Where:

- **Cfg_SR_MTU** is the MTU configured by the user for all SR tunnels within a given 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 (explained below).
- **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 IGP and passed to the SR module each time it changes due to an updated resolution of the node SID.

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

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

$$\text{SRTE\_LSP\_MTU} = \text{SR\_Tunnel\_MTU} - \text{numLabels} * 4$$

Where:

- **SR_Tunnel_MTU** is the MTU SR tunnel shortest path the SR-TE LSP is riding. The formula is as given above.
- **numLabels** is the number of labels found in the super NHLFE of the SR-TE LSP. Note that at LER, the super NHLFE is pointing to the SR tunnel NHLFE, which itself has a primary and a backup NHLFEs.

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 riding is updated.



Note: For the purpose of fragmentation of IP packets forwarded in GRT or in a VPRN over a SR shortest path tunnel or a SR-TE LSP, the IOM always deducts the worst case MTU (12 labels) from the outgoing interface MTU for the decision to fragment or not the packet. In this case, the above formula is not used.

2.8.9.2 Hash Label and Entropy Label Support

When the **hash-label** option is enabled in a service context, hash label is always inserted at the bottom of the stack as per RFC 6391.

The LSR adds the capability to hash up to a maximum of 16 labels in a stack. The LSR is able to hash on the IP headers when the payload below the label stack of maximum size of 16 is IPv4 or IPv6, including when a MAC header precedes it (**ethencap-ip** option).

The Entropy Label (EL) feature, as specified in RFC 6790, is introduced for RSVP, LDP, and BGP transport tunnels. It uses the Entropy Label Indicator (ELI) to indicate the presence of the entropy label in the label stack. The ELI, followed by the actual entropy label, is inserted immediately below the transport label for which entropy label feature is enabled. If multiple transport tunnels have the entropy label feature enabled, the ELI/EL is inserted below the lowest transport label in the stack.

The Entropy Label feature is not supported with an SR-TE LSP. If, however, a BGP LSP is tunneled over a SR-TE LSP and the entropy label is enabled on the BGP LSP, the LSR parses the label stack and detects the ELI below the BGP label.

The LSR hashing operates as follows.

- If the **lbl-only** hashing option is enabled, or if one of the other LSR hashing options is enabled but an IPv4 or IPv6 header is not detected below the bottom of the label stack, the LSR hashes on the EL only.
- If the **lbl-ip** option is enabled, the LSR hashes on the EL and the IP headers.
- If the **ip-only** or **eth-encap-ip** is enabled, the LSR hashes on the IP headers only.

See the [MPLS Entropy Label and Hash Label](#) section for more information about the hash label and Entropy Label features.

2.9 MPLS Service Usage

Nokia routers enable service providers to deliver virtual private networks (VPNs) and Internet access using Generic Routing Encapsulation (GRE) and/or MPLS tunnels, with Ethernet interfaces and/or SONET/SDH (on the 7750 SR and 7450 ESS) interfaces.


2.9.1 Service Distribution Paths

A service distribution path (SDP) acts as a logical way of directing traffic from one router to another through a uni-directional (one-way) service tunnel. The SDP terminates at the far-end router which directs packets to the correct service egress service access point (SAP) on that device. All services mapped to an SDP use the same transport encapsulation type defined for the SDP (either GRE or MPLS).

For information about service transport tunnels, refer to the Service Distribution Paths (SDPs) section in the OS Services Guide. They can support up to eight forwarding classes and can be used by multiple services. Multiple LSPs with the same destination can be used to load-balance traffic.

2.10 MPLS/RSVP Configuration Process Overview

[Figure 39](#) displays the process to configure MPLS and RSVP parameters.

Figure 39 MPLS and RSVP Configuration and Implementation Flow

al_0212

2.11 Configuration Notes

This section describes MPLS and RSVP caveats.

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

2.12 Configuring MPLS and RSVP with CLI

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

Topics in this section include:

- [MPLS Configuration Overview](#)
 - [LSPs](#)
 - [Paths](#)
 - [Router Interface](#)
 - [Choosing the Signaling Protocol](#)
- [Basic MPLS Configuration](#)
- [Common Configuration Tasks](#)
 - [Configuring MPLS Components](#)
 - [Configuring Global MPLS Parameters](#)
 - [Configuring an MPLS Interface](#)
 - [Configuring MPLS Paths](#)
 - [Configuring an MPLS LSP](#)
 - [Configuring Manual Bypass Tunnels](#)
- [Configuring and Operating SR-TE](#)
 - [SR-TE Configuration Prerequisites](#)
 - [SR-TE LSP Configuration Overview](#)
 - [Configuring Path Computation and Control for SR-TE LSP](#)
 - [Configuring SR-TE LSP Label Stack Size](#)
 - [Configuring Adjacency SID Parameters](#)
 - [Configuring PCC-controlled, PCE-computed, and PCE-controlled SR-TE LSPs](#)
- [Configuring RSVP Parameters](#)
 - [Configuring RSVP Message Pacing Parameters](#)
 - [Configuring Graceful Shutdown](#)
- [MPLS Configuration Management Tasks](#)
- [RSVP Configuration Management Tasks](#)

2.13 MPLS Configuration Overview

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

In order to implement MPLS, the following entities must be configured:

- [LSPs](#)
- [Paths](#)
- [Router Interface](#)

2.13.1 LSPs

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

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

2.13.2 Paths

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

2.13.3 Router Interface

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

2.13.4 Choosing the Signaling Protocol

In order to configure a static or a RSVP signaled LSP, you must enable MPLS on the router, which automatically enables RSVP and adds the system interface into both contexts. Any other network IP interface, other than loopbacks, added to MPLS is also automatically enabled in RSVP and becomes a TE link. When the interface is enabled in RSVP, the IGP instance will advertise the Traffic Engineering (TE) information for the link to other routers in the network in order to build their TE database and compute CSPF paths. Operators must enable the traffic-engineering option in the ISIS or OSPF instance for this. Operators can also configure under the RSVP context of the interface the RSVP protocol parameters for that interface.

If only static label switched paths are used in your configurations, operators must manually define the paths through the MPLS network. Label mappings and actions configured at each hop must be specified. Operators can disable RSVP on the interface if it is used only for incoming or outgoing static LSP label by shutting down the interface in the RSVP context. The latter causes IGP to withdraw the TE link from its advertisement which removes it from its local and neighbors TE database.

If dynamic LSP signaling is implemented in an operator's network then they must keep RSVP enabled on the interfaces they want to use for explicitly defined or CSPF calculated LSP path.

2.14 Basic MPLS Configuration

This section provides information to configure MPLS and configuration examples of common configuration tasks. To enable MPLS, you must configure at least one MPLS interface. The other MPLS configuration parameters are optional. This follow displays an example of an MPLS configuration.

```
ALA-1>config>router>if-attr# info
-----
admin-group "green" 15
admin-group "yellow" value 20
admin-group "red" value 25
-----
A:ALA-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 "secondary-path"
  no shutdown
exit
path "to-NYC"
  hop 1 10.10.10.104 strict
  no shutdown
exit
lsp "lsp-to-eastcoast"
  to 10.10.10.104
  from 10.10.10.103
  fast-reroute one-to-one
  exit
  primary "to-NYC"
  exit
  secondary "secondary-path"
  exit
  no shutdown
exit
static-lsp "StaticLabelPush"
  to 10.10.11.105
  push 60 nexthop 10.10.11.105
  no shutdown
exit
no shutdown
-----
A:ALA-1>config>router>mpls#

```

2.15 Common Configuration Tasks

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

The following protocols must be enabled on each participating router.

- MPLS
- RSVP (for RSVP-signaled MPLS only), which is automatically enabled when MPLS is enabled.

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

- An interface must be created in the **config>router>interface** context before it can be applied to MPLS.

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

2.15.1 Configuring MPLS Components

Use the MPLS and RSVP CLI syntax displayed below for:

- [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 Parameters](#)
- [Configuring RSVP Message Pacing Parameters](#)
- [Configuring Graceful Shutdown](#)

2.15.2 Configuring Global MPLS Parameters

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

To configure MPLS admin-group parameters, enter the following commands:

CLI Syntax:

```

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

The following displays an admin group configuration example:

```
ALA-1>config>router>if-attr# info
```

```

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

```

2.15.3 Configuring an MPLS Interface

Configure the **label-map** parameters if the interface is used in a static LSP. To configure an MPLS interface on a router, enter the following commands:

CLI Syntax:

```

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

```

The following displays an interface configuration example:

```

A:ALA-1>config>router>mpls# info
-----
...
        interface "to-104"
            admin-group "green"
            admin-group "red"
            admin-group "yellow"
            label-map 35
                swap 36 nexthop 10.10.10.91
                no shutdown
            exit
        exit
        no shutdown
...
-----
A:ALA-1>config>router>mpls#

```

2.15.4 Configuring MPLS Paths

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

Use the following 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 displays a path configuration example:

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

2.15.5 Configuring an MPLS LSP

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

The following displays an MPLS LSP configuration:

```
A:ALA-1>config>router>mpl# info
-----
...
lsp "lsp-to-eastcoast"
```

```

        to 192.168.200.41
        rsvp-resv-style ff
        cspf
        include "red"
        exclude "green"
        adspec
        fast-reroute one-to-one
        exit
        primary "to-NYC"
            hop-limit 10
        exit
        secondary "secondary-path"
            bandwidth 50000
        exit
        no shutdown
    exit
no shutdown
-----
A:ALA-1>config>router>mpls#

```

2.15.5.1 Configuring a Static LSP

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

Use the following CLI syntax to configure a static LSP:

CLI Syntax:

```

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

```

The following displays a static LSP configuration example:


```

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

```

2.15.6 Configuring Manual Bypass Tunnels

Consider the following network setup.



The user first configures the option to disable the dynamic bypass tunnels on node B if required. The CLI for this configuration is:

```
config>router>mpls>dynamic-bypass [disable | enable]
```

By default, dynamic bypass tunnels are enabled.

Next, the user configures an LSP on node B, such as B-E-F-C to be used only as bypass. The user specifies each hop in the path, for example, the bypass LSP has a strict path.

Note that including the bypass-only keyword disables the following options under the LSP configuration:

- bandwidth
- fast-reroute
- secondary

The following LSP configuration options are allowed:

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

The following example displays a bypass tunnel configuration:

```

A:ALA-48>config>router>mpls>path# info
-----
...
    path "BEFC"
        hop 10 10.10.10.11 strict
        hop 20 10.10.10.12 strict
  
```

```

        hop 30 10.10.10.13 strict
        no shutdown
    exit

    lsp "bypass-BC"
    to 10.10.10.15
    primary "BEFC"
    exit
    no shutdown
...
-----
A:ALA-48>config>router>mpls>path#

```

Next, the user configures an LSP from A to D and indicates fast-reroute bypass protection by selecting facility as the FRR method (**config>router>mpls>lsp>fast-reroute facility**). If the LSP goes through B, and bypass is requested, and the next hop is C, and there is a manually configured bypass-only tunnel from B to C, excluding link BC, then node B uses that.

2.16 Configuring and Operating SR-TE

This section provides information on the configuration and operation of the Segment Routing with Traffic Engineering (SR-TE) LSP.

- [SR-TE Configuration Prerequisites](#)
- [SR-TE LSP Configuration Overview](#)
- [Configuring Path Computation and Control for SR-TE LSP](#)
- [Configuring SR-TE LSP Label Stack Size](#)
- [Configuring Adjacency SID Parameters](#)
- [Configuring PCC-controlled, PCE-computed, and PCE-controlled SR-TE LSPs](#)

2.16.1 SR-TE Configuration Prerequisites

To configure SR-TE, the user must first configure 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:

```

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-alternate 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
  no shutdown
exit
no shutdown
exit
```

Configure an 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
```

2.16.2 SR-TE LSP Configuration Overview

An SR-TE LSP can be configured as a label switched path (LSP) using the existing CLI command hierarchy under the MPLS context and specifying the new **sr-te** LSP type.

CLI Syntax:

```
config>router>mpls>lsp lsp-name | mpls-tp src-tunnel-num
| sr-te
```

As for an RSVP LSP, the user can configure a primary path.

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
```

2.16.3 Configuring Path Computation and Control for SR-TE LSP

Use the following syntax to configure the path computation requests only (PCE-computed) or both path computation requests and path updates (PCE-controlled) to PCE for a specific LSP:

CLI Syntax:

```
config>router>mpls>lsp>pce-computation
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}
```

2.16.3.1 Configuring Path Profile and Group for PCC-Initiated and PCE-Computed/Controlled LSP

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 PCE, the user must configure the PCE path profile ID and path group ID the LSP belongs to. These parameters are passed transparently by PCC to PCE and are thus 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 PCE 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 section of this guide.

2.16.4 Configuring SR-TE LSP Label Stack Size

Use the following syntax to configure the maximum number of labels which the ingress LER can push for a given 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, service, and other labels when packets are forwarded in a given 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 could not 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 1-10 labels with a default value of 6.

2.16.5 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 of the interval during which 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
no entropy-label-capability
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-alternate** option in ISIS 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 a SR-TE LSP. In that case, use the following syntax to disable protection for all adjacency SIDs formed over a given 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
```

2.16.6 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 138.120.48.226
no shutdown
exit
no shutdown
```

The following example shows the configuration of a PCC-controlled SR-TE LSP that is not reported to PCE:

Example:

```
lsp "to-SanFrancisco" sr-te
  to 38.120.48.211
  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 PCE:

Example:

```
lsp "to-SanFrancisco" sr-te
  to 38.120.48.211
  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 PCE:

Example:

```
lsp "to-SanFrancisco" sr-te
  to 38.120.48.211
  cspf
  pce-computation
  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 38.120.48.224
  cspf
  pce-computation
  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 38.120.48.224
cspf
max-sr-labels 8 additional-frr-labels 1
pce-computation
pce-report enable
pce-control
path-profile 10 path-group 2
primary "empty"
    bandwidth 15
exit
no shutdown
exit

```

2.17 Configuring RSVP Parameters

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

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

The following example displays an RSVP configuration example:

```

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

```

2.17.1 Configuring RSVP Message Pacing Parameters

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

Use the following CLI syntax to configure RSVP parameters:

CLI Syntax:

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

The following example displays a RSVP message pacing configuration example:

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

2.17.2 Configuring Graceful Shutdown

TE graceful shutdown can be enabled on a specific interface using the **config>router>rsvp>if>graceful-shutdown** command. This interface is referred to as the maintenance interface.

Graceful shutdown can be disabled by executing the **no** form of the command at the RSVP interface level or at the RSVP level. In this case, the user configured TE parameters of the maintenance links are restored and the maintenance node floods them.

2.18 MPLS Configuration Management Tasks

This section discusses the following MPLS configuration management tasks:

- [Deleting MPLS](#)
- [Modifying MPLS Parameters](#)
- [Modifying MPLS Path Parameters](#)
- [Modifying MPLS Static LSP Parameters](#)
- [Deleting an MPLS Interface](#)

2.18.1 Deleting MPLS

NOTE: In order to remove the MPLS instance, MPLS must be disabled (shutdown) and all SDP bindings to LSPs removed. If MPLS is not shutdown first, when the **no mpls** command is executed, a warning message on the console displays indicating that MPLS is still administratively up.

When MPLS is shut down, the **no mpls** command deletes the protocol instance and removes all configuration parameters for the MPLS instance. To disable MPLS, use the **shutdown** command.

To remove MPLS on a router, enter the following command:

CLI Syntax: `config>router# no mpls`

2.18.2 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.

2.18.3 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 displays a MPLS LSP configuration example. Refer to the LSP configuration in [Configuring an MPLS LSP](#).

```
A:ALA-1>>config>router>mpls>lsp# info
-----
      shutdown
      to 10.10.10.104
      from 10.10.10.103
      rsvp-resv-style ff
      include "red"
      exclude "green"
      fast-reroute one-to-one
      exit
      primary "to-NYC"
        hop-limit 50
      exit
      secondary "secondary-path"
      exit
-----
A:ALA-1>config>router>mpls#
```

2.18.4 Modifying MPLS Path Parameters

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

The following displays a path configuration example. Refer to [Configuring MPLS Paths](#).

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

2.18.5 Modifying MPLS Static LSP Parameters

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

The following displays a static LSP configuration example. Refer to the static LSP configuration on [2.15.5.1](#).

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

2.18.6 Deleting an MPLS Interface

In order to delete an interface from the MPLS configuration, the interface must be shut down first.

Use the following CLI syntax to delete an interface from the MPLS configuration:

CLI Syntax:

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

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

2.19 RSVP Configuration Management Tasks

This section discusses the following RSVP configuration management tasks:

- [Modifying RSVP Parameters](#)
- [Modifying RSVP Message Pacing Parameters](#)
- [Deleting an Interface from RSVP](#)

2.19.1 Modifying RSVP Parameters

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

The **no rsvp** command deletes this RSVP protocol instance and removes all configuration parameters for this RSVP instance.

The **shutdown** command suspends the execution and maintains the existing configuration.

The following example displays a modified RSVP configuration example:

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

2.19.2 Modifying RSVP Message Pacing Parameters

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

The following example displays command usage to modify RSVP parameters:

The following example displays a modified RSVP message pacing configuration example. Refer to [Configuring RSVP Message Pacing Parameters](#).

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

2.19.3 Deleting an Interface from RSVP

Interfaces cannot be deleted directly from the RSVP configuration. An interface must have been configured in the MPLS context, which enables it automatically in the RSVP context. The interface must first be deleted from the MPLS context. This removes the association from RSVP.

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

2.20 MPLS/RSVP Configuration Command Reference

2.20.1 Command Hierarchies

- [MPLS Commands](#)
- [MPLS-TP Commands](#)
- [LSP Commands](#)
- [lsp-bfd Commands](#)
- [MPLS Path Commands](#)
- [RSVP Commands](#)

2.20.1.1 MPLS Commands

```

config
  — router
    — [no] mpls
      — [no] admin-group-frr
      — auto-bandwidth-multipliers sample-multiplier number1 adjust-multiplier number2
      — no auto-bandwidth-multipliers
      — auto-lsp lsp-template template-name {policy peer-prefix-policy [peer-prefix-policy...(upto 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}
      — exponential-backoff-retry
      — [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] **p2p-template-lsp** **rsvp-session-name** *SessionNameString*
 - sender** *sender-address*
 - **accounting-policy** *policy-id*
 - **no accounting-policy**
 - [no] **collect-stats**
 - [no] **max-stats**
 - [no] **shutdown**
 - [no] **p2mp-template-lsp** **rsvp-session-name** *SessionNameString*
 - sender** *sender-address*
 - **accounting-policy** *policy-id*
 - **no accounting-policy**
 - [no] **collect-stats**
 - [no] **max-stats**
 - [no] **shutdown**
 - [no] **interface** *ip-int-name*
 - [no] **admin-group** *group-name* [*group-name...*(up to 5 max)]
 - **no admin-group**
 - [no] **label-map** *in-label*
 - [no] **pop**
 - [no] **shutdown**
 - **swap** {*out-label* | **implicit-null-label**} **nexthop** *ip-addr*
 - **no swap** {*out-label* | **implicit-null-label**}
 - [no] **mpls-tp-mep**
 - [no] **ais-enable**
 - **if-num** *if-num*
 - **no if-num**
 - **if-num-validation** {**enable** | **disable**}
 - [no] **shutdown**
 - [no] **srlg-group** *group-name* [*group-name...*(up to 5 max)]
 - **no srlg-group**
 - **te-metric** *metric*
 - **no te-metric**
 - **least-fill-min-thd** *percent*
 - **no least-fill-min-thd**
 - **least-fill-reoptim-thd** *percent*
 - **no least-fill-reoptim-thd**
 - **lsp-init-retry-timeout** *seconds*
 - **no lsp-init-retry-timeout**
 - [no] **logger-event-bundling**
 - **max-bypass-associations** *integer*
 - **no max-bypass-associations**
 - [no] **mbb-prefer-current-hops**
 - **p2p-active-path-fast-retry** *seconds* [1..10] *seconds*
 - **no p2p-active-path-fast-retry**
 - **p2mp-s2l-fast-retry** *seconds* [1..10] *seconds*
 - **no p2mp-s2l-fast-retry**
 - **preemption-timer** *seconds*
 - **no preemption-timer**
 - **p2mp-resignal-timer** *minutes*
 - **no p2mp-resignal-timer**
 - **pce-report sr-te** {**enable** | **disable**}
 - **resignal-timer** *minutes*
 - **no resignal-timer**
 - [no] **retry-on-igp-overload**

- **secondary-fast-retry-timer** *seconds*
- **no secondary-fast-retry-timer**
- [no] **shortcut-local-ttl-propagate**
- [no] **shortcut-transit-ttl-propagate**
- [no] **shutdown**
- [no] **srlg-database**
 - [no] **router-id** *ip*
 - [no] **interface** *ip-addr srlg-group group-name [group-name..(up to 5 max)]*
 - [no] **shutdown**
- [no] **srlg-frr** [strict]
- **srlg-group**
- **srlg-group**
- [no] **static-lsp** *lsp-name*
 - **metric** *metric*
 - **no metric**
 - **push** {*label* | **implicit-null-label**} **nexthop** *ip-address*
 - **no push** {*out-label* | **implicit-null-label**}
 - [no] **shutdown**
 - **to***ip-address*
- **static-lsp-fast-retry** *seconds*
- **no static-lsp-fast-retry**
- **user-srlg-db** [enable | disable]

2.20.1.2 MPLS-TP Commands

- ```

config
 — router
 — [no] mpls
 — [no] mpls-tp
 — global-id global-id
 — no global-id
 — node-id node-id
 — no node-id
 — [no] oam-template name
 — hold-time-down timer
 — no hold-time-down
 — hold-time-up timer
 — no hold-time-up
 — bfd-template name
 — no bfd-template
 — protection-template name
 — no protection-template
 — [no] revertive
 — wait-to-restore interval
 — no wait-to-restore
 — rapid-psc-timer interval
 — no rapid-psc-timer
 — slow-psc-timer interval

```

- **no slow-psc-timer**
- **[no] shutdown**
- **tp-tunnel-id-range** *start-id end-id*
- **no tp-tunnel-id-range**
- **transit-path** *path-name*
- **no transit-path**
  - **[no] forward-path**
  - **in-label** *in-label out-label out-label out-link if-name [next-hop next-hop]*
  - **no in-label**
  - **[no] mep**
    - **dsmap** *if-num*
    - **no dsmap**
  - **path-id** {*lsp-num lsp-num | working-path | protect-path [src-global-id src-global-id] src-node-id src-node-id src-tunnel-num src-tunnel-num [dest-global-id dest-global-id] dest-node-id dest-node-id [dest-tunnel-num dest-tunnel-num]*}
  - **no path-id**
  - **[no] reverse-path**
  - **[no] shutdown**

### 2.20.1.3 LSP Commands

- ```

config
  — router
    — [no] mpls
      — [no] lsp lsp-name [bypass-only | p2mp-lsp | mpls-tp src-tunnel-num | sr-te]
        — [no] adaptive
        — [no] adspec
        — [no] auto-bandwidth
          — adjust-down percent [bw mbps]
          — no adjust-down
          — adjust-up percent [bw mbps]
          — no adjust-up
          — max-bandwidth mbps
          — no max-bandwidth
          — min-bandwidth mbps
          — no min-bandwidth
          — [no] monitor-bandwidth
          — multipliers sample-multiplier num1 adjust-multiplier num2
          — no multipliers
          — overflow-limit number threshold percent [bw mbps]
          — no overflow-limit
          — underflow-limit number threshold percent [bw mbps]
          — no underflow-limit
        — bfd
          — bfd-enable
          — no bfd-enable
          — bfd-template name
          — no bfd-template

```


- **lsp-ping-interval** *seconds*
- **no lsp-ping-interval**
- **[no] bgp-shortcut**
- **bgp-transport-tunnel** *include | exclude*
- **[no] class-forwarding**
 - **fc** {**be** | **l2** | **af** | **l1** | **h2** | **ef** | **h1** | **nc**}
 - **no fc** [{**be** | **l2** | **af** | **l1** | **h2** | **ef** | **h1** | **nc**}]
 - **[no] default-lsp**
- **class-type** *ct-number*
- **no class-type**
- **[no] cspf** [**use-te-metric**]
- **dest-global-id** *dest-global-id*
- **no dest-global-id**
- **dest-tunnel-number** *dest-tunnel-number*
- **no dest-tunnel-number**
- **[no] egress-statistics**
 - **accounting-policy** *policy-id*
 - **no accounting-policy**
 - **[no] collect-stats**
 - **[no] shutdown**
- **entropy-label** {**force-disable** | **enable** | **inherit**}
- **[no] exclude** *group-name* [*group-name...*(up to 5 max)]
- **[no] exclude-node** *ip-address*
- **fast-reroute** *frr-method*
- **no fast-reroute**
 - **[no] propagate-admin-group**
 - **hop-limit** *number*
 - **no hop-limit**
 - **[no] node-protect**
- **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] ingress-statistics**
- **ldp-over-rsvp** [**include** | **exclude**]
- **[no] least-fill**
- **[no] ldp-over-rsvp** [**include** | **exclude**]
- **load-balancing-weight** *integer* (32-bit)
- **no load-balancing-weight**
- **main-ct-retry-limit** *number*
- **no main-ct-retry-limit**
- **max-sr-labels** *label-stack-size* [**additional-frr-labels** *labels*]
- **no max-sr-labels**
- **[no] metric** *metric*
- **p2mp-id** *id*
- **path-profile** *profile-id* [**path-group** *group-id*]
- **no path-profile**
- **[no] pce-computation**
- **[no] pce-control**
- **pce-report** {**enable** | **disable** | **inherit**}
- **[no] primary** *path-name*
 - **[no] adaptive**

-
- **backup-class-type** *ct-number*
 - **no backup-class-type**
 - **bandwidth** *rate-in-mpbs*
 - **no bandwidth**
 - **bfd**
 - **bfd-enable**
 - **no bfd-enable**
 - **bfd-template** *name*
 - **no bfd-template**
 - **lsp-ping-interval** *seconds*
 - **no lsp-ping-interval**
 - **class-type** *ct-number*
 - **no class-type**
 - **[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)]
 - **priority** *setup-priority hold-priority*
 - **no priority**
 - **[no] record**
 - **[no] record-label**
 - **[no] shutdown**
 - **[no] primary-p2mp-instance** *instance-name*
 - **[no] adaptive**
 - **bandwidth** *rate-in-mbps*
 - **no bandwidth**
 - **[no] exclude** *group-name* [*group-name...*(up to 5 max)]
 - **[no] hop-limit**
 - **hop-limit** *number*
 - **no hop-limit**
 - **[no] include** *group-name* [*group-name...*(up to 5 max)]
 - **[no] record**
 - **[no] record-label**
 - **[no] s2l-path** *path-name* **to** *ip-address*
 - **[no] shutdown**
 - **[no] shutdown**
 - **[no] propagate-admin-group**
 - **[no] protect-tp-path**
 - **bfd-enable** [**cc** | **cc-cv**]
 - **no bfd-enable**
 - **[no] mep**
 - **[no] bfd-trap-suppression**
 - **dsmap** *if-num*
 - **no dsmap**
 - **in-label** *in-label*
 - **no in-label**
 - **lsp-num** *lsp-num*
 - **no lsp-num**
 - **[no] mep**
 - **oam-template** *name*
 - **no oam-template**
 - **out-label** *out-label* **out-link** *if-name* [**next-hop** *ip-address*]
 - **no out-label**
 - **protection-template** *name*

```

    — no protection-template
    — [no] shutdown
— retry-limit number
— no retry-limit
— retry-timer seconds
— no retry-timer
— [no] revert-timer timer-value
— rsvp-resv-style [se | ff]
— [no] secondary path-name
    — [no] adaptive
    — bandwidth rate-in-mbps
    — no bandwidth
    — class-type ct-number
    — no class-type
    — [no] exclude group-name [group-name...(up to 5 max)]
    — hop-limit number
    — no hop-limit
    — [no] include group-name [group-name...(up to 5 max)]
    — [no] path-preference preference-number
    — priority setup-priority hold-priority
    — no priority
    — [no] record
    — [no] record-label
    — [no] shutdown
    — [no] srlg
    — [no] standby
— [no] shutdown
— to [ip-address | node-id [a.b.c.d. | 1...4,294,967,295]]
— vprn-auto-bind [include | exclude]
— [no] working-tp-path
    — bfd-enable [cc | cc_cv]
    — no bfd-enable
    — in-label in-label
    — no in-label
    — lsp-num lsp-num
    — no lsp-num
    — [no] mep
        — [no] bfd-trap-suppression
        — dsmap if-num
        — no dsmap
    — oam-template name
    — no oam-template
    — out-label out-label out-link if-name [next-hop ip-address]
    — no out-label
    — [no] shutdown
— lsp-template template-name [p2mp | one-hop-p2p | mesh-p2p]
— no lsp-template template-name
    — [no] adspec
    — [no] auto-bandwidth
        — adjust-down percent [bw mbps]
        — no adjust-down
        — adjust-up percent [bw mbps]
        — no adjust-up
        — fc fc-name sampling-weight sampling-weight

```

```

— no fc
— max-bandwidth mbps
— no max-bandwidth
— min-bandwidth mbps
— no min-bandwidth
— [no] monitor-bandwidth
— multipliers sample-multiplier num1 adjust-multiplier num2
— no multipliers
— overflow-limit number threshold percent [bw mbps]
— no overflow-limit
— underflow-limit number threshold percent [bw mbps]
— no underflow-limit
— [no] bandwidth rate-in-mbps
— bfd
  — bfd-enable
  — no bfd-enable
  — bfd-template name
  — no bfd-template
  — lsp-ping-interval seconds
  — no lsp-ping-interval
— [no] class-forwarding
  — fc {be | l2 | af | l1 | h2 | ef | h1 | nc}
  — no fc [{be | l2 | af | l1 | h2 | ef | h1 | nc}]
  — [no] default-lsp
— [no] cspf [use-te-metric]
— [no] default-path path-name
— [no] egress-statistics
  — accounting-policy policy-id
  — no accounting-policy
  — [no] collect-stats
— entropy-label {force-disable | enable | inherit}
— [no] entropy-label
— [no] exclude-node ip-address
— fast-reroute frr-method
— no fast-reroute
  — [no] propagate-admin-group
  — hop-limit number
  — no hop-limit
  — [no] node-protect
— 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)]
— ldp-over-rsvp [include | exclude]
— [no] least-fill
— load-balancing-weight integer (32-bit)
— no load-balancing-weight
— [no] metric metric
— [no] propagate-admin-group
— [no] record
— [no] record-label
— retry-limit number

```

- **no retry-limit**
- **retry-timer** *seconds*
- **no retry-timer**
- **vprn-auto-bind** [*include* | *exclude*]

2.20.1.4 Lsp-bfd Commands

```

config
  — router
    — lsp-bfd
    — no lsp-bfd
      — bfd-sessions max-limit
      — no bfd-sessions

```

2.20.1.5 MPLS Path Commands

```

config
  — router
    — [no] mpls
      — [no] path path-name
        — hop hop-index ip-address {strict | loose}
        — no hop hop-index
        — [no] shutdown
      — [no] static-lsp lsp-name
        — metric metric
        — no metric
        — push label nexthop ip-address
        — no push out-label
        — to ip-addr
        — [no] shutdown

```

2.20.1.6 RSVP Commands

```

config
  — router
    — [no] rsvp
      — diffserv-te [mam | rdm]
      — no diffserv-te
        — class-type-bw ct0 %-link-bandwidth ct1 %-link-bandwidth ct2 %-link-bandwidth ct3 %-link-bandwidth ct4 %-link-bandwidth ct5 %-link-bandwidth ct6 %-link-bandwidth ct7 %-link-bandwidth
        — no class-type-bw
        — fc fc-name class-type ct-number
        — no fc fc-name
        — te-class te-class-number class-type ct-number priority priority

```

- **no te-class** *te-class-number*
- **[no] entropy-label-capability**
- **gr-helper-time max-recovery** *recovery-interval [1..1800] seconds max-restart restart-interval*
- **no gr-helper-time**
- **[no] graceful-shutdown**
- **[no] implicit-null-label**
- **[no] interface** *ip-int-name*
 - **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
 - **no authentication-key**
 - **auth-keychain** *name*
 - **no auth-keychain**
 - **[no] bfd-enable**
 - **class-type-bw** **ct0** %link-bandwidth **ct1** %link-bandwidth **ct2** %link-bandwidth **ct3** %link-bandwidth **ct4** %link-bandwidth **ct5** %link-bandwidth **ct6** %link-bandwidth **ct7** %link-bandwidth
 - **no class-type-bw**
 - **gr-helper** [**enable** | **disable**]
 - **[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**
 - **te-up-threshold** *threshold-level [threshold-level,...(up to 16 max)]*
 - **no te-up-threshold**
 - **te-down-threshold** *threshold-level [threshold-level,...(up to 16 max)]*
 - **no te-down-threshold**
- **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**]
- **p2p-merge-point-abort-timer** [*1..65535*] *seconds*
- **no p2p-merge-point-abort-timer**
- **p2mp-merge-point-abort-timer** [*1..65535*] *seconds*
- **no p2mp-merge-point-abort-timer**
- **preemption-timer** *seconds*
- **no preemption-timer**
- **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**

- [no] **te-threshold-update**
 - [no] **on-cac-failure**
 - **update-timer** *seconds*
 - **no update-timer**
- **te-up-threshold** *threshold-level* [*threshold-level...*(up to 16 max)]
- **no te-up-threshold**
- **te-down-threshold** *threshold-level* [*threshold-level...*(up to 16 max)]
- **no te-down-threshold**

2.20.2 Command Descriptions

- [MPLS Commands](#)
- [RSVP Commands](#)

2.20.2.1 MPLS Commands

2.20.2.1.1 Generic Commands

shutdown

Syntax	[no] shutdown
Context	config>router>mpls config>router>mpls>interface config>router>mpls>ingr-stats>p2mp-template-lsp config>router>mpls>ingr-stats>p2p-template-lsp config>router>mpls>lsp>primary config>router>mpls>lsp>primary-p2mp-instance config>router>mpls>lsp>secondary
Description	This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. The config>router>mpls>ingr-stats>p2mp-template-lsp>shutdown command is supported on the 7750 SR, 7950 XRS, and with VPLS only on the 7450 ESS. The config>router>mpls>lsp>primary-p2mp-instance>shutdown is not supported on the 7450 ESS. MPLS is not enabled by default and must be explicitly enabled (no shutdown).

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

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

Default no shutdown

2.20.2.1.2 MPLS Commands

mpls

Syntax [no] mpls

Context config>router

Description This command enables the context to configure MPLS parameters. MPLS is not enabled by default and must be explicitly enabled (**no shutdown**). The **shutdown** command administratively disables MPLS.

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

MPLS must be **shutdown** before the MPLS instance can be deleted. All SDP bindings to LSPs must be removed before the MPLS instance can be deleted..

If MPLS is not shutdown, when the **no mpls** command is executed, a warning message on the console displays indicating that MPLS is still administratively up.

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

Context config>router>mpls>ingr-stats>lsp
config>router>mpls>ingr-stats>p2mp-template-lsp
config>router>mpls>ingr-stats>p2p-template-lsp
config>router>mpls>lsp>egr-stats
config>router>mpls>lsp>ingr-stats
config>router>mpls>lsp>template>egr-stats

Description This command associates an accounting policy to the MPLS instance.

The config>router>mpls>ingr-stats>p2mp-template-lsp>accounting-policy command is supported on the 7750 SR, 7950 XRS, and with VPLS only on the 7450 ESS.

An accounting policy must be defined before it can be associated else an error message is generated.

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

Default	none
Parameters	<i>acct-policy-id</i> — Enter the accounting <i>policy-id</i> as configured in the config>log>accounting-policy context.
Values	1 to 99

collect-stats

Syntax	[no] collect-stats
Context	config>router>mpls>ingr-stats config>router>mpls>ingr-stats>p2mp-template-lsp config>router>mpls>ingr-stats>p2p-template-lsp config>router>mpls>lsp>egr-stats config>router>mpls>lsp-template>egr-stats
Description	This command enables accounting and statistical data collection. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file. The config>router>mpls>ingr-stats>p2mp-template-lsp>collect-stats command is supported on the 7750 SR, 7950 XRS, and with VPLS only on the 7450 ESS. When 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 billing file. If a subsequent collect-stats command is issued then the counters written to the billing file include all the traffic while the no collect-stats command was in effect.
Default	collect-stats

max-stats

Syntax	[no] max-stats
Context	config>router>mpls>ingr-stats config>router>mpls>ingr-stats>p2mp-template-lsp config>router>mpls>ingr-stats>p2p-template-lsp config>router>mpls>lsp>egr-stats config>router>mpls>lsp-template>egr-stats
Description	This command enables accounting and statistical data collection. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

The `config>router>mpls>ingr-stats>p2mp-template-lsp>max-stats` command is supported on the 7750 SR, 7950 XRS, and with VPLS only on the 7450 ESS.

When the **no max-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 billing file. If a subsequent **max-stats** command is issued then the counters written to the billing file include all the traffic while the **no max-stats** command was in effect.

Default max-stats

dynamic-bypass

Syntax **dynamic-bypass** [enable | disable]
no dynamic-bypass

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}

Context config>router>mpls

Description This command configures the use of entropy labels for MPLS.

The entropy label and entropy label indicator (ELI) require 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 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, regardless of the value set for **entropy-label** at the egress node, the ingress LER will consider the entropy label and ELI in the label stack while sending the information to the TTM and NHLFE. The application using the LSP will not insert an entropy label and ELI in the label stack unless the far-end signals ELC and the application is configured to insert an entropy label.

When the **entropy-label** command value changes at either the MPLS level or the LSP level, the new operational value will not take effect until the LSP is re-signaled. A **shutdown/no shutdown** of the LSP is required to enable the new value.

The user can use the **clear** command or bounce MPLS itself (**shutdown/no shutdown**) to force the new value to take effect for a large numbers of LSPs.

Default	entropy-label enable
Parameters	rsvp-te — Indicates that the entropy-label command applies to RSVP LSPs force-disable — Indicates that 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 and ELI in the label stack. enable — Indicates that 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 will always be accounted for.

exponential-backoff-retry

Syntax	exponential-backoff-retry no exponential-backoff-retry
Context	config>router>mpls
Description	This command enables the use of an exponential back-off timer when re-trying an LSP. When an LSP path establishment attempt fails, the path is put into retry procedures and a new attempt will be performed at the expiry of the user-configurable retry timer (config>router>mpls>lsp>retry-timer). By default, the retry time is constant for every attempt. The exponential back-off timer procedures will double the value of the user configured retry timer value at every failure of the attempt to adjust to the potential network congestion that caused the failure. An LSP establishment fails if no Resv message was received and the Path message retry timer expired or a PathErr message was received before the timer expired.

admin-group-frr

Syntax	[no] admin-group-frr
Context	config>router>mpls
Description	This command enables the use of the 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, then the PLR will read 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, then it just 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 among those that are already in use.

If none of the manual or dynamic bypass LSP 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 that includes or excludes the specified admin-group IDs.

If the user changes the configuration of the above command, it will not have any effect on existing bypass associations. The change will only apply 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 frr-admin-group

frr-object

Syntax [no] frr-object

Context config>router>mpls

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

Default frr-object — The value is by default 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 holds before programming its data plane and declaring the LSP up to the service module. This occurs anytime the ingress node brings up an LSP path or switches traffic from a working path to another working path of the same LSP.

The **no** form of the command reverts the hold-timer to the default value.

Default	1
Parameters	<i>seconds</i> — Specifies the time (in s), for which the ingress node holds before programming its data plane and declaring the LSP up to the service module.
Values	0 to 10

ingress-statistics

Syntax	ingress-statistics
Context	config>router>mpls>lsp
Description	This command provides the context for the user to enter the LSP names for the purpose of enabling ingress data path statistics at the terminating node of the LSP, for example, egress LER.
Default	none

ingress-statistics

Syntax	ingress-statistics
Context	config>router>mpls
Description	This command provides the context for the user to enable ingress-statistics on an MPLS-TP LSP.
Default	none

least-fill-min-thd

Syntax	least-fill-min-thd <i>percent</i> no least-fill-min-thd
Context	config>router>mpls
Description	This parameter is used in the least-fill path selection process. When comparing the percentage of least available link bandwidth across the sorted paths, whenever two percentages differ by less than the value configured as the least-fill-min-thresh, CSPF will consider them equal and will apply a random number generator to select the path among these paths The no form of the command resets this parameter to its default value.
Default	5

Parameters *percentage* — 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 method. During a timer-based re-signaling of an LSP path which has the least-fill option enabled, CSPF will first update the least-available bandwidth figure 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, it will compare the least-available bandwidth figures of the two paths and if the difference exceeds the user configured optimization threshold, MPLS will generate 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 re-signaling of the LSP path since the timer-based re-signaling will not re-signal the path in this case. MPLS will generate a path update trap at the first MBB event which results in the re-signaling of the LSP path. This should clear the eligibility status of the path at the SNMP device.

The **no** form of this command resets this parameter to its default value.

Default 10

Parameters *percentage* — Specifies the least fill reoptimization threshold value as a percentage.
Values 1 to 100%

lsp

Syntax [**no**] **lsp** *lsp-name* **sender** *sender-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 operator at the ingress LER. It must not contain the special character “.” 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 operator must execute the **no shutdown** for this command to effectively enable statistics.

The same set of counters is updated for packets received over any path of this LSP and over the lifetime of the LSP. In 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.

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. A consequence of this is that when the operator enables statistics collection on an RSVP LSP which is also used for tunneling LDP FECs with the LDP over RSVP feature, then statistics will be collected on the RSVP LSP only. There will be no statistics collected for an LDP FEC tunneled over this RSVP LSP and also egressing on the same node regardless if the operator enabled statistics collection on this FEC. When, the operator disables statistics collection on the RSVP LSP, then statistics collection, if enabled, will be performed on a tunneled LDP FEC.

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

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	none
Parameters	sender-address <i>ip-address</i> — A string of 15 characters representing the IP address of the ingress LER for the LSP. <i>lsp-name</i> — A string of up to 32 characters identifying the LSP name as configured at the ingress LER.

p2p-template-lsp

Syntax	[no] p2p-template-lsp rsvp-session-name <i>SessionNameString</i> sender <i>sender-address</i>
Context	config>router>mpls>ingress-stats
Description	This command configures an ingress statistics context matching on the RSVP session name for a RSVP P2P auto-LSP at the egress LER.

When the ingress LER signals the path of a template based **one-hop-p2p** or **mesh-p2p auto-lsp**, it includes the name of the LSP and that of the path in the Session Name field of the Session Attribute object in the Path message. The encoding is as follows:

Session Name: *lsp-name::path-name*, where *lsp-name* component is encoded as follows:

P2MP LSP via user configuration for L3 multicast in global routing instance:
“LspNameFromConfig”

- P2MP LSP as I-PMSI or S-PMSI in L3 mVPN: *templateName-SvcId-mTTmIndex*
- P2MP LSP as I-PMSI in VPLS/B-VPLS: *templateName-SvcId-mTTmIndex*.

The ingress statistics CLI configuration allows the user to match either on the exact name of the P2P auto-LSP or on a context that matches on the template name and the destination of the LSP at the ingress LER.

When the matching is performed on a context, the user must enter the RSVP session name string in the format “templateName-svcId” to include the LSP template name as well as the mVPN VPLS/B-VPLS service ID as configured at the ingress LER. In this case, one or more P2MP LSP instances signaled by the same ingress LER could be associated with the ingress statistics configuration. In this case, the user is provided with CLI parameter max-stats to limit the maximum number of stat indices which can be assigned to this context. If the context matches more than this value, the additional request for stat indices from this context will be rejected.

Note the following rules when configuring an ingress statistics context based on template matching:

- **max-stats**, once allocated, can be increased but not decreased unless the entire ingress statistics context matching a template name is deleted.
- In order to delete ingress statistics context matching a template name, a shutdown is required.
- An accounting policy cannot be configured or de-configured until the ingress statistics context matching a template name is shut down.
- After deleting an accounting policy from an ingress statistics context matching a template name, the policy is not removed from the log until a **no shut** is performed on the ingress statistics context.

If there are no stat indices available at the time the session of the P2P LSP matching a template context is signaled and the session state installed by the egress LER, no stats are allocated to the session.

Furthermore, the assignment of stat indices to the LSP names that match the context will also be not deterministic. The latter is due to the fact that a stat index is assigned and released following the dynamics of the LSP creation or deletion by the ingress LER. For example, a multicast stream crosses the rate threshold and is moved to a newly signaled S-PMSI dedicated to this stream. Later on, the same stream crosses the threshold downwards and is moved back to the shared I-PMSI and the P2MP LSP corresponding to the S-PMSI is deleted by the ingress LER.

The **no** form deletes the ingress statistics context matching on the RSVP session name.

- Parameters**
- rsvp-session-name** *SessionNameString* — Specifies the name of the RSVP P2MP session in the format of “templateName-svcId”. This field can hold up to 64 characters.
- sender** *sender-address* — Specifies a string of 15 characters representing the IP address of the ingress LER for the LSP.

p2mp-template-lsp

Syntax	[no] p2mp-template-lsp rsvp-session-name <i>SessionNameString</i> sender <i>sender-address</i>
Context	config>router>mpls>ingress-stats
Description	This command configures an ingress statistics context matching on the RSVP session name for a RSVP P2MP LSP at the egress LER.

This command is supported on the 7750 SR, 7950 XRS, and with VPLS only on the 7450 ESS.

When the ingress LER signals the path of the S2L sub-LSP, it includes the name of the LSP and that of the path in the Session Name field of the Session Attribute object in the Path message. The encoding is as follows:

Session Name: <lsp-name::path-name>, where lsp-name component is encoded as follows:

- P2MP LSP via user configuration for L3 multicast in global routing instance: "LspNameFromConfig"
- P2MP LSP as I-PMSI or S-PMSI in L3 mVPN: templateName-SvcId-mTTmIndex
- P2MP LSP as I-PMSI in VPLS/B-VPLS: templateName-SvcId-mTTmIndex

The ingress statistics CLI configuration allows the user to match either on the exact name of the P2MP LSP as configured at the ingress LER or on a context that matches on the template name and the service-id as configured at the ingress LER.

When the matching is performed on a context, the user must enter the RSVP session name string in the format "templateName-svcId" to include the LSP template name as well as the mVPN VPLS/B-VPLS service ID as configured at the ingress LER. In this case, one or more P2MP LSP instances signaled by the same ingress LER could be associated with the ingress statistics configuration and the user is provided with CLI parameter max-stats to limit the maximum number of stat indices that can be assigned to this context. If the context matches more than this value, the additional request for stat indices from this context will be rejected. A background tasks monitors the ingress statistics templates which have one or more matching LSP instances with unassigned stat index and assigns one to them as they are freed.

Note the following rules when configuring an ingress statistics context based on template matching:

- max-stats, once allocated, can be increased but not decreased unless the entire ingress statistics context matching a template name is deleted.
- In order to delete ingress statistics context matching a template name, a shutdown is required.
- An accounting policy cannot be configured or de-configured until the ingress statistics context matching a template name is shut down.
- After deleting an accounting policy from an ingress statistics context matching a template name, the policy is not removed from the log until a "no shut" is performed on the ingress statistics context.

If there are no stat indices available at the time the session of the P2MP LSP matching a template context is signaled and the session state installed by the egress LER, no stats are allocated to the session.

Furthermore, the assignment of stat indices to the LSP names that match the context will also be not deterministic. The latter is due to the fact that a stat index is assigned and released following the dynamics of the LSP creation or deletion by the ingress LER. For example, a multicast stream crosses the rate threshold and is moved to a newly signaled S-PMSI dedicated to this stream. Later on, the same steam crosses the threshold downwards and is moved back to the shared I-PMSI and the P2MP LSP corresponding to the S-PMSI is deleted by the ingress LER.

The **no** form deletes the ingress statistics context matching on the RSVP session name.

- Parameters**
- rsvp-session-name** *SessionNameString* — Specifies the name of the RSVP P2MP session in the format of “templateName-svcld”. This field can hold up to 64 characters.
 - sender** *sender-address* — Specifies a string of 15 characters representing the IP address of the ingress LER for the LSP.

logger-event-bundling

- Syntax** **[no] logger-event-bundling**
- Context** config>router>mpls
- Description** This feature merges two of the most commonly generated MPLS traps, vRtrMplsXCCreate and vRtrMplsXCDelete, which can be generated at both LER and LSR into a new specific trap vRtrMplsSessionsModified. In addition, this feature will perform bundling of traps of multiple RSVP sessions, i.e., LSPs, into this new specific trap.
- The intent is to provide a tool for the user to minimize trap generation in an MPLS network. Note that the MPLS trap throttling will not be applied to this new trap.
- The **no** version of this command disables the merging and bundling of the above MPLS traps.

lsp-template

- Syntax** **lsp-template** template-name [**p2mp** | **one-hop-p2p** | **mesh-p2p**]
 no lsp-template *template-name*
- Context** config>router>mpls
- Description** This command creates a template construct that can be referenced by client application where dynamic LSP creation is required. The LSP template type **p2mp**, **one-hop-p2p**, or **mesh-p2p** is mandatory.

The **no** form of command deletes LSP template. LSP template cannot be deleted if a client application is using it.

Parameters *lsp-template-name* — Specifies the name of the LSP template. Any LSP template name and LSP name must not be the same.

p2mp | one-hop-p2p | mesh-p2p — Identifies the type of the LSP this template will signal.

The **p2mp** option is supported on the 7750 SR, 7950 XRS, and with VPLS only on the 7450 ESS.

lsp-init-retry-timeout

Syntax **lsp-init-retry-timeout** *seconds*
no lsp-init-retry-timeout

Context config>router>mpls

Description This command configures the initial LSP path retry-timer.

The new LSP path initial retry-timer is used instead of the retry-timer to abort the retry cycle when no RESV is received. The retry-timer will govern exclusively the time between two retry cycles and to handle retrying of an LSP path in a failure case with PATH errors or RESVTear.

The intent is that the user can now control how many refreshes of the pending PATH state can be performed before starting a new retry-cycle with a new LSP-id. This is all done without affecting the ability to react faster to failures of the LSP path, which will continue to be governed by the retry-timer.

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

Parameters *seconds* — Specifies the value (in s), used as the fast retry timer for a secondary path.

Values 10to600

Default 30

lsp-template

Syntax **lsp-template** *template-name* [**p2mp** | **one-hop-p2p** | **mesh-p2p**]
no lsp-template *template-name*

Context config>router>mpls

Description This command creates a template construct that can be referenced by client application where dynamic LSP creation is required. The LSP template type **p2mp**, **one-hop-p2p**, or **mesh-p2p** is mandatory.

The **no** form of command deletes LSP template. LSP template cannot be deleted if a client application is using it.

- Parameters**
- lsp-template-name* — Specifies the name to identify LSP template. An LSP template name and LSP name must not be the same.
 - p2mp | one-hop-p2p | mesh-p2p** — Identifies the type of the LSP this template will signal.

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 FRR object at the ingress.

When this command is executed, the admin-group constraints configured in the context of the P2P LSP primary path, or the ones 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 thus propagates these constraints to the downstream nodes during the signaling of the LSP to allow them to include the admin-group constraints in the selection of the FRR backup LSP for protecting the LSP primary path.

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

Note that the same admin-group constraints can be copied into the Session Attribute object. 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 regardless if the hop is strict or loose. These are governed strictly by the command:

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

In other words, the user may decide to copy the primary path admin-group constraints into the FAST_REROUTE object only, or into the Session Attribute object only, or into both. Note, however, that the PLR rules for processing the admin-group constraints can make use of either of the two object admin-group constraints.

This feature is supported with the following LSP types and in both intra-area and inter-area TE where applicable:

- Primary path of a RSVP P2P LSP.
- S2L path of an RSVP P2MP LSP instance

- LSP template for an S2L path of an RSVP P2MP LSP instance.

The **no** form of this command disables the signaling of administrative group constraints in the FRR object.

Default no propagate-admin-group

max-bypass-associations

Syntax	max-bypass-associations <i>integer</i> no max-bypass-associations
Context	config>router>mpls
Description	This command allows the user to set a maximum number of LSP primary path associations with each manual or dynamic bypass LSP that is created in the system. By default, a Point of Local Repair (PLR) node will associate a maximum of 1000 primary LSP paths with a given bypass before using the next available manual bypass or signaling a new dynamic bypass. Note that a new bypass LSP may need to be signaled if the constraint of a given primary LSP path is not met by an existing bypass LSP even if the max-bypass-associations for this bypass LSP has not been reached. The no form of the command re-instates the default value of this parameter.
Default	no max-bypass-associations
Parameters	<i>integer</i> — Configures the number of LSP primary path associations
Values	1 to 131,072

mbb-prefer-current-hops

Syntax	[no] mbb-prefer-current-hops
Context	config>router>mpls
Description	This command implements a new option in the CSPF path computation during a Make-Before-Break (MBB) procedure of an RSVP LSP. When MPLS performs an MBB for the primary or secondary path of a P2P LSP, or the S2L path of a P2MP LSP, and the new mbb-prefer-current-hops option is enabled in MPLS context, CSPF will select a path, among equal-cost candidate paths, with the most overlapping links with the current path. Normally, CSPF selects the path randomly.

The procedures of the new MBB CSPF path selection apply to LSP without the least-fill option enabled. If the least-fill rule results in a different path, the LSP path will be moved though. Users can still favor stability over least-fill condition by applying a larger value to the parameter **least-fill-min-thd** under the MPLS context such that a path will only be moved when the difference of the least-available bandwidth becomes significant enough between the most used links in the equal cost paths. If that difference is not significant enough, CSPF will select the path with the most overlapping links instead of selecting a path randomly.

The procedures when the new **mbb-prefer-current-hops** option is enabled apply to all MBB types. Thus, it applies to the auto-bandwidth MBB, the configuration change MBB, the soft pre-emption MBB, the TE graceful shutdown MBB, the delayed retry MBB (for SRLG secondary LSP path), the path change MBB, the timer resignal MBB, and the manual resignal MBB.

During the FRR global revertive MBB, CSPF selects a random link among the ones available between the PLR node and the Merge Point node, including the failed link if it has restored in the meantime. These links cannot be checked for overlap with the current path.

The TE graceful shutdown MBB will still avoid the link or node that is in maintenance and the soft pre-emption MBB will still avoid the link that is overbooked.

For an inter-area LSP, this feature applies to the subset of the path from the ingress LER to the exit ABR.

The procedures of this feature are not applied to a zero bandwidth CSFP LSP, including an auto-bandwidth CSPF LSP while its operational bandwidth is zero, and to a non-CSPF LSP.

pce-report

Syntax	pce-report sr-te {enable disable} no pce-report sr-te
Context	config>router>mpls
Description	This command configures the reporting mode to a PCE of all SR-TE LSPs in a PCC. 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/disable PCE reporting for all SR-TE LSPs for the purpose of LSP database synchronization. This configuration will be inherited by all LSPs of a given type. The PCC will report both CSPF and non-CSPF LSP. The default value is disabled (no pce-report). This default value is meant to control the introduction of PCE into an existing network and let the operator decide if all LSP types need to be reported. The LSP level pce-report command overrides the global configuration for the reporting of LSP to PCE. The default value is to inherit the global MPLS level value. The inherit value returns the LSP to inherit the global configuration for that LSP type.

Note that 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.

Default	no pce-report
Parameters	enable disable — Specifies to enable or disable configuring PCE report types for MPLS.

resignal-timer

Syntax	resignal-timer <i>minutes</i> 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, the software waits before attempting to resignal the LSPs. When the resignal timer expires, if the new computed path for an LSP has a better metric than the current recorded hop list, an attempt is made to resignal that LSP using the make-before-break mechanism. If the attempt to resignal an LSP fails, the LSP will continue to use the existing path and a resignal will be attempted the next time the timer expires. The no form of the command disables timer-based LSP resignaling.
Default	no resignal-timer
Parameters	<i>minutes</i> — The time the software waits before attempting to resignal the LSPs. Values 30 to 10080

retry-on-igp-overload

Syntax	[no] retry-on-igp-overload
Context	config>router>mpls
Description	This command allows for the global configuration of the handling in the ingress LER of the LSP paths which transit an LSR that advertised the IS-IS overload bit. By default, MPLS re-optimizes using make-before-break (MBB) the transit paths away from the node in an IS-IS overload state only at the time a manual or timer-based re-signal is performed for the LSP paths. MPLS will not act immediately on the receipt of the IS-IS overload bit. When this command is enabled, MPLS in the ingress LER immediately tears down and re-signals all LSP paths away from a transit LSR node which advertised the IS-IS overload bit.

LSP paths that terminate on the node that advertised the IS-IS overload bit are not acted on whether this command is enabled or disabled.

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

secondary-fast-retry-timer

Syntax	secondary-fast-retry-timer <i>seconds</i> no secondary-fast-retry-timer
Context	config>router>mpls
Description	This command specifies the value used as the fast retry timer for a secondary path. If the first attempt to set up a secondary path fails due to a path error, the fast retry timer will be started for the secondary path so that the path can be retried sooner. If the next attempt also fails, further retries for the path will use the configured value for LSP retry timer. If retry-timer for the LSP is configured to be less than the MPLS secondary-fast-retry-timer, all retries for the secondary path will use the LSP retry-timer. The no form of the command reverts to the default.
Default	no secondary-fast-retry-timer
Parameters	<i>seconds</i> — specifies the value (in s), used as the fast retry timer for a secondary path
Values	1 to 10

shortcut-local-ttl-propagate

Syntax	[no] shortcut-local-ttl-propagate
Context	config>router>mpls
Description	This command enables or disables TTL propagation over an LSP shortcut for local packets.

shortcut-transit-ttl-propagate

Syntax	[no] shortcut-transit-ttl-propagate
Context	config>router>mpls
Description	This command enables or disables TTL propagation over an LSP shortcut for transit packets..

srlg-frr

Syntax	srlg-frr [strict] no srlg-frr
Context	config>router>mpls
Description	This command enables the use of the Shared Risk Loss Group (SRLG) constraint in the computation of FRR bypass or detour to be associated with any primary LSP path on this system.

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

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

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

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

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

Once the bypass or detour path is setup and is operationally UP, any subsequent changes to the SRLG group membership of an interface the bypass/detour path is using would not be considered by the MPLS/RSVP task at the PLR until the next opportunity the association with the primary LSP path is re-checked. The association is re-checked if the bypass path is re-optimized. Detour paths are not re-optimized and are re-signaled if the primary path is down.

Enabling or disabling srlg-frr only takes effect at the next opportunity the LSP paths are resignaled. The user can wait for the resignal timer to expire or can cause the paths to be resignaled immediately by executing at the ingress LER the **tools perform router mpls resignal** command. Note that in order to force the dynamic bypass LSP to be resignaled using the SRLG constraint of the primary paths it is associated with, it is recommend to first

disable dynamic bypass LSPs on the system using the “configure router mpls dynamicbypass” command, then manually resignal the LSP paths using the above tools perform command finally re-enable dynamic bypass LSPs on the system. Before performing this procedure, the user must ensure that no dynamic bypass LSP on the node is active to avoid causing the primary LSP path to go down.

An RSVP interface can belong to a maximum of 64 SRLG groups. The user configures the SRLG groups using the command **config>router>mpls>srlg-group**. The user configures the SRLG groups an RSVP interface belongs to 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 the name of the SRLG group within a virtual router instance.
	Values no slr-frr (default) srlg-frr (non-strict) srlg-frr strict (strict)

srlg-group

Syntax	[no] srlg-group <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no srlg-group
Context	config>router>if>if-attribute config>service>ies>if>if-attribute config>service>vprn>if>if-attribute config>router>mpls>interface
Description	This command configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface.

An interface can belong to up to 64 SRLG groups. However, each single operation of the srlg-group command allows a maximum of five (5) groups to be specified at a time. Once an SRLG group is bound to one or more interface, its value cannot be changed until all bindings are removed.

The configured SRLG membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/ areas.

It should be noted that only the SRLGs bound to an MPLS interface are advertised in TE link TLVs and sub-TLVs when the traffic-engineering option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

The no form of this command deletes one or more of the SRLG memberships of an interface.

The user can also delete all memberships of an interface by not specifying a group name.

Default	no srlg-group
Parameters	<i>group-name</i> — Specifies the name of the group, up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain.

user-srlg-db

Syntax	user-srlg-db [enable disable]
Context	config>router>mpls
Description	This command enables the use of CSPF by the user SRLG database. When the MPLS module makes a request to CSPF for the computation of an SRLG secondary path, CSPF will query the local SRLG and compute a path after pruning links that are members of the SRLG IDs of the associated primary path. When MPLS makes a request to CSPF for an FRR bypass or detour path to associate with the primary path, CSPF queries the user SRLG database and computes a path after pruning links that are members of the SRLG IDs of the PLR outgoing interface. If an interface was not entered into the user SRLG database, it is assumed that it does not have any SRLG membership. CSPF will not query the TE database for IGP advertised interface SRLG information. The disable keyword disables the use of the user SRLG database. CSPF will then resume queries into the TE database for SRLG membership information. The user SRLG database is maintained.
Default	user-srlg-db disable

srlg-database

Syntax	[no] srlg-database
Context	config>router>mpls
Description	This command provides the context for the user to enter manually the link members of SRLG groups for the entire network at any node that needs to signal LSP paths (for example, a head-end node). The no form of the command deletes the entire SRLG database. CSPF will assume all interfaces have no SRLG membership association if the database was not disabled with the command config>router>mpls>user-srlg-db disable.

router-id

Syntax	[no] router-id <i>ip</i>

Context	config>router>mpls>srlg-database
Description	This command provides the context for the user to manually enter the link members of SRLG groups for a specific router in the network. The user must also use this command to enter the local interface SRLG membership into the user SRLG database. Use by CSPF of all interface SRLG membership information of a specific router ID may be temporarily disabled by shutting down the node. If this occurs, CSPF will assume these interfaces have no SRLG membership association. The no form of this command will delete all interface entries under the router ID.
Parameters	<i>ip-address</i> — Specifies the router ID for this system. This must be the router ID configured under the base router instance, the base OSPF instance or the base IS-IS instance.

interface

Syntax	interface <i>ip-address srlg-group group-name</i> [<i>group-name...</i> (up to 5 max)] no interface <i>ip-address [srlg-group group-name...</i> (up to 5 max)]
Context	config>router>mpls>srlg-database>router-id
Description	This command allows the operator to manually enter the SRLG membership information for any link in the network, including links on this node, into the user SRLG database. An interface can be associated with up to 5 SRLG groups for each execution of this command. The operator can associate an interface with up to 64 SRLG groups by executing the command multiple times. CSPF will not use entered SRLG membership if an interface is not validated as part of a router ID in the routing table. The no form of the command deletes a specific interface entry in this user SRLG database. The group-name must already exist in the config>router>mpls>srlg-group context.
Default	none
Parameters	<i>ip-int-name</i> — The name of the network IP interface. An interface name cannot be in the form of an IP address. <i>srlg-group group-name</i> — Specifies the SRLG group name. Up to 1024 group names can be defined in the config>router>mpls context. The SRLG group names must be identical across all routers in a single domain.

load-balancing-weight

Syntax	load-balancing-weight <i>integer</i> no load-balancing-weight

Context	config>router>mpls>lsp
Description	This command assigns a weight to an MPLS LSP for use in the weighted load-balancing, or weighted ECMP, over MPLS feature.
Default	none
Parameters	<i>value</i> — 32-bit integer representing the weight of the LSP.
Values	0 to 4294967295

MPLS Interface Commands

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>router>mpls
Description	This command specifies MPLS protocol support on an IP interface. No MPLS commands are executed on an IP interface where MPLS is not enabled. An MPLS interface must be explicitly enabled (no shutdown). The no form of this command deletes all MPLS commands such as label-map which are defined under the interface. The MPLS interface must be shutdown first in order to delete the interface definition. If the interface is not shutdown, the no interface <i>ip-int-name</i> command does nothing except issue a warning message on the console indicating that the interface is administratively up.
Default	shutdown
Parameters	<i>ip-int-name</i> — The name of the network IP interface. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. Values 1 to 32 alphanumeric characters.

admin-group

Syntax	[no] admin-group <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no admin-group
Context	config>router>if>if-attribute config>service>ies>if>if-attribute config>service>vpn>if>if-attribute config>router>mpls>interface

Description	This command configures the admin group membership of an interface. The user can apply admin groups to an IES, VPRN, network IP, or MPLS interface. Each single operation of the admin-group command allows a maximum of five (5) groups to be specified at a time. However, a maximum of 32 groups can be added to a given interface through multiple operations. Once an admin group is bound to one or more interface, its value cannot be changed until all bindings are removed. The configured admin-group membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas. It should be noted that only the admin groups bound to an MPLS interface are advertised in TE link TLVs and sub-TLVs when the traffic-engineering option is enabled in IS-IS or OSPF. The IES and VPRN interfaces do not have their attributes advertised in TE TLVs. The no form of this command deletes one or more of the admin-group memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.
Default	no admin-group
Parameters	<i>group-name</i> — Specifies the name of the group with up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain.

auto-bandwidth-multipliers

Syntax	auto-bandwidth-multipliers sample-multiplier <i>number1</i> adjust-multiplier <i>number2</i> no auto-bandwidth-multipliers
Context	config>router>mpls
Description	This command specifies the number of collection intervals in the adjust interval.
Parameters	sample-multiplier <i>number1</i> — Specifies the multiplier for collection intervals in a sample interval. Values 1 to 511 Default 1 adjust-multiplier <i>number2</i> — Specifies the number of collection intervals in the adjust interval. Values 1 to 16383 Default 288

auto-lsp

Syntax	auto-lsp lsp-template <i>template-name</i> {policy <i>peer-prefix-policy</i> [peer-prefix-policy...(upto 5 max)] one-hop} no auto-lsp lsp-template <i>template-name</i>

Context	config>router>mpls
Description	This command enables the automatic creation of an RSVP 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 auto-LSP of type mesh. The user can associate multiple templates with same or different peer prefix policies. Each application of an LSP template with a given prefix in the prefix list will result 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. This feature does not support the automatic signaling of a secondary path for an LSP. If the user requires the signaling of multiple LSPs to the same destination node, s/he must apply a separate LSP template to the same or different prefix list that contains the same destination node. Each instantiated LSP will have a unique LSP-id and a unique tunnel-ID. This feature also does not support the signaling of a non-CSPF LSP. The selection of the no cspf option in the LSP template is thus blocked. Up to five (5) peer prefix policies can be associated with a given LSP template at all times. Each time the user executes the above command with the same or different prefix policy associations, or the user changes a prefix policy associated with an LSP template, the system re-evaluates the prefix policy. The outcome of the re-evaluation will tell MPLS if an existing LSP needs to be torn down or if a new LSP needs to be signaled to a destination address that is already in the TE database. If a /32 prefix is added to (removed from) or if a prefix range is expanded (shrunk) in a prefix list associated with a LSP template, the same prefix policy re-evaluation described above is performed. The user must perform a no shutdown of the template before it takes effect. Once a template is in use, the user must shutdown the template before effecting any changes to the parameters except for those LSP parameters for which the change can be handled with the Make-Before-Break (MBB) procedures. These parameters are bandwidth and enabling fast-reroute with or without the hop-limit or node-protect options. For all other parameters, the user shuts down the template and once it is added, removed or modified, the existing instances of the LSP using this template are torn down and re-signaled. The trigger to signal the LSP is when the router with a router-id the matching a prefix in the prefix list appears in the Traffic Engineering database. The signaled LSP is installed in the Tunnel Table Manager (TTM) and is available to applications such as LDP-over-RSVP, resolution of BGP label routes, resolution of BGP, IGP, and static routes. It is, however, not available to be used as a provisioned SDP for explicit binding or auto-binding by services. Except for the MBB limitations to the configuration parameter change in the LSP template, MBB procedures for manual and timer based re-signaling of the LSP, for TE Graceful Shutdown and for soft pre-emption are supported. The one-to-one option under fast-reroute, the LSP Diff-Serv class-type and backup-class-type parameters are not supported. If diffserv-te is enabled under RSVP, the auto-created LSP will still be signaled but with the default LSP class type.

If the **one-hop** option is specified instead of a prefix list, this command enables the automatic signaling of one-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. Although the provisioning model and CLI syntax differ from that of a mesh LSP only by the absence of a prefix list, the actual behavior is quite different. When the above command is executed, the TE database will keep track of each TE link that comes up to a directly connected IGP neighbor which 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. Thus, the **auto-lsp** command with the **one-hop** option will result in one or more LSPs signaled to the neighboring router.

An auto-created mesh or one-hop LSP can have egress statistics collected at the ingress LER by adding the **egress-statistics** node configuration into the LSP template. The user can also have **ingress statistics** collected at the egress LER using the same ingress-statistics node in CLI used with a provisioned LSP. The user must specify the full LSP name as signaled by the ingress LER in the RSVP session name field of the Session Attribute object in the received Path message.

The **no** form of this command deletes all LSP signaled using the specified template and prefix policy. When the **one-hop** option is used, it deletes all one-hop LSPs signaled using the specified template to all directly connected neighbors.

Parameters **lsp-template** *template-name* — Specifies an LSP template name up to 32 characters in length.

policy *peer-prefix-policy* — Specifies an peer prefix policy name up to 32 characters in length.

bypass-resignal-timer

Syntax **bypass-resignal-timer** *minutes*
 no bypass-resignal-timer

Context config>router>mpls

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

When this command is enabled, MPLS makes a request to CSPF for the best path for each dynamic bypass LSP originated on this node. The constraints of the first associated LSP primary path and which 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 that LSP primary path even if the latter is torn down.

If CSPF returns no path or returns a new path that is equal in terms of 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 will signal it.

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

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

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

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

The periodic bypass re-optimization 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.

Furthermore, PSBs that have not been moved by the dynamic or manual re-optimization 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.

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

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

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

Default	no bypass-resignal timer. The periodic global re-optimization of dynamic bypass LSP paths is disabled.
Parameters	<i>minutes</i> — Specifies the time, in minutes, MPLS waits before attempting to re-signal 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>rsvp
Description	This command enables the option to do CSPF calculations until the next loose hop or the final destination of LSP on LSR. On receiving a PATH message on LSR and processing of all local hops in the received ERO, if the next hop is loose, then the LSR node will first do a CSPF calculation until the next loose hop. On successful completion of CSPF calculation, ERO in PATH message is modified to include newly calculated intermediate hops and propagate it forward to the next hop. This allows setting up inter-area LSPs based on ERO expansion method. NOTE: The LSP may fail to set up if this option is enabled on an LSR that is not an area border router and receives a PATH message without proper next loose hop in ERO. The 'cspf-on-loose-hop' configuration is allowed to change dynamically and applied to new LSP setup after change.
Default	no cspf-on-loose-hop

srlg-group

Syntax	[no] srlg-group <i>group-name</i> [<i>group-name...</i> (up to 5 max)]
Context	config>router>mpls>interface
Description	This command defines the association of RSVP interface to an SRLG group. An interface can belong to up to 64 SRLG groups. However, each single operation of the srlg-group command allows a maximum of 5 groups to be specified at a time. The no form of this command deletes the association of the interface to the SRLG group.
Default	none
Parameters	<i>group-name</i> — Specifies the name of the SRLG group within a virtual router instance up to 32 characters.

te-metric

Syntax	te-metric <i>value</i> 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 IS-IS TE metric is encoded as sub-TLV 18 as part of the extended IS reachability TLV. The metric value is encoded as a 24-bit unsigned integer. The OSPF TE metric is encoded as a sub-TLV Type 5 in the Link TLV. The metric value is encoded as a 32-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 which 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>cspf>use-te-metric**.

Note that 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

The value of the IGP metric is advertised in the TE metric sub-TLV by IS-IS and OSPF.

Parameters *value* — Specifies the metric value.

Values 1 to 16777215

node-id-in-rro

Syntax [no] node-id-in-rro [include | exclude]

Context config>router>rsvp>

Description This command enables the option to include node-id sub-object in RRO. Node-ID sub-object propagation is required to provide fast reroute protection for LSP that spans across multiple area domains.

If this option is disabled, then node-id is not included in RRO object.

Default node-id-in-rro exclude

p2p-merge-point-abort-timer

Syntax p2p-merge-point-abort-timer [1.. 65535] seconds
no p2p-merge-point-abort-timer

Context config>router>rsvp>

Description	This command configures a timer to abort Merge-Point (MP) node procedures for a P2P LSP path. When a value higher than zero is configured for this timer, it will enter into effect anytime this node activates Merge-Point procedures for one or more P2P LSP paths. As soon an ingress interface goes operationally down, the Merge-Point node starts the abort timer. Upon expiry of the timer, MPLS will clean up all P2P LSP paths which ILM is on the failed interface and which have not already received a Path refresh over the bypass LSP.
Default	0 (disabled)

p2mp-merge-point-abort-timer

Syntax	p2mp-merge-point-abort-timer [1.. 65535] <i>seconds</i> no p2mp-merge-point-abort-timer
Context	config>router>rsvp
Description	This command specifies a configurable timer to abort Merge-Point (MP) node procedures for an S2L of a P2MP LSP instance. When a value higher than zero is configured for this timer, it will enter into effect anytime this node activates Merge-Point procedures for one or more P2MP LSP S2L paths. As soon an ingress interface goes operationally down, the Merge-Point node starts the abort timer. Upon expiry of the timer, MPLS will clean up all P2MP LSP S2L paths which ILM is on the failed interface and which have not already received a Path refresh over the bypass LSP.
Default	0 (disabled)

p2p-active-path-fast-retry

Syntax	p2p-active-path-fast-retry <i>seconds</i> no p2p-active-path-fast-retry
Context	config>router>mpls
Description	This command configures a global parameter to allow the user to apply a shorter retry timer for the first try after an active LSP path went down due to a local failure or the receipt of a RESVTear. This timer is used only in the first try. Subsequent retries will continue to be governed by the existing LSP level retry-timer.
Default	0 (disabled)
Parameters	<i>seconds</i> — Specifies the length of time for retry timer, in seconds Values 1 to 10 seconds

p2mp-s2l-fast-retry

Syntax	p2mp-s2-fast-retry <i>seconds</i>

no p2mp-s2l-fast-retry

Context	config>router>mpls config>router>rsvp
Description	This command configures a global parameter to allow the user to apply a shorter retry timer for the first try after an active LSP path went down due to a local failure or the receipt of a RESVTear. This timer is used only in the first try. Subsequent retries will continue to be governed by the existing LSP level retry-timer. The config>router>mpls>p2mp-s2l-fast-retry command is supported on the 7750 SR, 7950 XRS, and with VPLS only on the 7450 ESS.
Default	0 (disabled)
Parameters	seconds — Specifies the length of time for retry timer, in seconds
Values	1 to 10 seconds

preemption-timer

Syntax	preemption-timer seconds no preemption-timer
Context	config>router>rsvp
Description	This parameter configures the time in seconds a node holds to a reservation for which it triggered the soft pre-emption procedure. The pre-empting node starts a separate preemption timer for each pre-empted LSP path. While this timer is on, the node should continue to refresh the Path and Resv for the pre-empted LSP paths. When the preemption timer expires, the node tears down the reservation if the head-end node has not already done so. A value of zero means the LSP should be pre-empted immediately; hard pre-empted. The no form of this command reverts to the default value.
Default	300
Parameters	seconds — Specifies the time (in s), of the preemption timer.
Values	0 to 1800 seconds

label-map

Syntax	[no] label-map in-label
Context	config>router>mpls>interface

Description	This command is used on transit routers when a static LSP is defined. The static LSP on the ingress router is initiated using the config router mpls static-lsp <i>lsp-name</i> command. An <i>in-label</i> can be associated with either a pop or a swap action, but not both. If both actions are specified, the last action specified takes effect. The no form of this command deletes the static LSP configuration associated with the <i>in-label</i>.
Parameters	<i>in-label</i> — 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 <i>in-label</i>.
Default	none

shutdown

Syntax	[no] shutdown
Context	config>router>mpls>if>label-map
Description	This command disables the label map definition. This drops all packets that match the specified <i>in-label</i> specified in the label-map <i>in-label</i> command. The no form of this command administratively enables the defined label map action.
Default	no shutdown

swap

Syntax	swap { <i>out-label</i> implicit-null-label } nexthop <i>ip-address</i> no swap { <i>out-label</i> implicit-null-label }
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	none
Parameters	<i>implicit-null-label</i> — Specifies the use of the implicit label value for the outgoing label of the swap operation. <i>out-label</i> — Specifies the label value to be swapped with the in-label. Label values 16 through 1,048,575 are defined as follows: Label values 16 through 31 are reserved. Label values 32 through 1,023 are available for static assignment. Label values 1,024 through 2,047 are reserved for future use. Label values 2,048 through 18,431 are statically assigned for services. Label values 28,672 through 131,071 are dynamically assigned for both MPLS and services. Label values 131,072 through 1,048,575 are reserved for future use. Values 16 to 1048575 nexthop <i>ip-address</i> — The IP address to forward to. If an ARP entry for the next hop exists, then the static LSP will be marked operational. If ARP entry does not exist, software will set the operational status of the static LSP to down and continue to ARP for the configured nexthop. Software will continuously try to ARP for the configured nexthop at a fixed interval.

mpls-tp-mep

Syntax	[no] mpls-tp-mep
Context	config>router>mpls>interface
Description	This command enables the context for a section layer MEP for MPLS-TP on an MPLS interface.
Default	none

ais-enable

Syntax	[no] ais-enable
Context	config>router>mpls>if>mpls-tp-mep
Description	This command enables MPLS-TP AIS insertion for the forward and reverse directions of all MPLS-TP transit paths using the MPLS interface. This causes the generation of AIS packets in the forward or reverse directions of a path if a fault is detected on the applicable underlying interface for the ingress of the path direction. The no form of this command disables AIS insertion.

Default no ais-enable

if-num

Syntax **if-num** *if-num*
 no if-num

Context config>router>mpls>if>mpls-tp-mep

Description This command configures the MPLS-TP interface number for the MPLS interface. This is a 32-bit unsigned integer that is node-wide unique.

Parameters *if-num* — This is a 32-bit value that is unique to the node.

Values 1 to 4294967295

if-num-validation

Syntax **if-num-validation** {**enable** | **disable**}
 no if-num-validation

Context config>router>mpls>if>mpls-tp-mep

Description The if-num-validation command is used to enable or disable validation of the if-num in LSP Trace packet against the locally configured if-num for the interface over which the LSP Trace packet was received at the egress LER. This is because some 3rd-party implementations may not perform interface validation for unnumbered MPLS-TP interfaces and instead set the if-num in the dsmmap TLV to 0. If the value is 'enabled', the node performs the validation of the ingress and egress if-nums received in the LSP echo request messages that ingress on this MPLS-interface. It validates that the message arrives on the interface as identified by the ingress if-num, and is forwarded on the interface as identified by the egress if-num.

 If the value is 'disabled', no validation is performed for the ingress and egress if-nums received in the LSP echo request messages that ingress on this MPLS-interface."

Default enable

Parameters **enable** — Enables interface number validation.
 disable — Disables interface number validation.

MPLS-TP Commands

mpls-tp

Syntax	[no] mpls-tp
Context	config>router>mpls
Description	Generic MPLS-TP parameters and MPLS-TP trabsit paths are configured under this context. If a user configures no mpls, normally the entire mpls configuration is deleted. However, in the case of mpls-tp, a check is made that there is no other mpls-tp configuration (e.g., services or LSPs using mpls-tp on the node). The mpls-tp context cannot be deleted if MPLS-TP LSPs or SDPs exist on the system. A shutdown of mpls-tp will bring down all MPLS-TP LSPs on the system.
Default	no mpls-tp

tp-tunnel-id-range

Syntax	tp-tunnel-id-range <i>start-id end-id</i> no tp-tunnel-id-range
Context	config>router>mpls>mpls-tp
Description	This command configures the range of MPLS tunnel IDs reserved for MPLS-TP LSPs. The maximum difference between the start-id and end-id is 4K. The tunnel ID referred to here is the RSVP-TE tunnel ID. This maps to the MPLS-TP Tunnel Number. There are some cases where the dynamic LSPs may have caused fragmentation to the number space such that contiguous range [<i>end-id</i> – <i>start-id</i>] is not available. In these cases, the command will fail. There are no default values for the <i>start-id</i> and <i>end-id</i> of the tunnel id range, and they must be configured to enable MPLS-TP.
Default	no tunnel-id-range
Parameters	<i>start-id</i> — Specifies the start ID. Values 1 to 61440 <i>end-id</i> — Specifies the end ID. Values 1 to 61440

oam-template

Syntax	[no] oam-template <i>name</i>
Context	config>router>mpls>mpls-tp
Description	This command creates or edits an OAM template. Generally applicable proactive OAM parameters are configured using templates. The top-level template is the OAM template. Generic MPLS-TP OAM and fault management parameters are configured in the OAM Template. Proactive CC/CV uses BFD and parameters such as Tx/Rx timer intervals, multiplier and other session/fault management parameters specific to BFD are configured using a BFD Template, which is referenced from the OAM template.
Default	no oam-template
Parameters	<i>name</i> — Specifies a text string name for the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes. Named OAM templates are referenced from the MPLS-TP path MEP configuration.

hold-time-down

Syntax	hold-time-down <i>timer</i> no hold-time-down
Context	config>router>mpls>mpls-tp>oam-template
Description	This command configures the hold-down dampening timer. It is equivalent to a hold-off timer.
Default	no hold-time-down
Parameters	<i>interval</i> — Specifies the hold-down dampening timer interval. Values 0 to 5000 deciseconds in 10 ms increments

hold-time-up

Syntax	hold-time-up <i>timer</i> no hold-time-up
Context	config>router>mpls>mpls-tp>oam-template
Description	This command configures the hold-up dampening timer. This can be used to provide additional dampening to the state of proactive CC BFD sessions.
Default	no hold-time-up

Parameters	<i>interval</i> — Specifies the hold-up dampening timer interval.
Values	0 to 500 deciseconds, in 100 ms increments
Default	2 s

bfd-template

Syntax	bfd-template <i>name</i> no bfd-template
Context	config>router>mpls>mpls-tp>oam-template
Description	This command configures a named BFD template to be referenced by an OAM template.
Default	no bfd-template
Parameters	<i>name</i> — Specifies the BFD template name as a text string up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

protection-template

Syntax	protection-template <i>name</i> no protection-template
Context	config>router>mpls>mpls-tp
Description	Protection templates are used to define generally applicable protection parameters for MPLS-TP tunnels. Only linear protection is supported, and so the application of a named template to an MPLS-TP LSP implies that linear protection is used. A protection template is applied under the MEP context of the protect-path of an MPLS-TP LSP. The protection-template command creates or edits a named protection template.
Default	no protection-template
Parameters	<i>name</i> — Specifies the protection template name as a text string of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

revertive

Syntax	[no] revertive
Context	config>router>mpls>mpls-tp>protection-template
Description	This command configured revertive behavior for MPLS-TP linear protection. The protect-tp-path MEP must be in the shutdown state for of the MPLS-TP LSPs referencing this protection template in order to change the revertive parameter.

Default revertive

wait-to-restore

Syntax **wait-to-restore** *interval*
 no wait-to-restore

Context config>router>mpls>mpls-tp>protection-template

Description This command configures the WTR timer. It determines how long to wait until the active path of an MPLS-TP LSP is restored to the working path following the clearing of a defect on the working path. It is applicable to revertive mode, only.

Default no wait-to-restore

Parameters *interval* — Specifies the WTR timer interval.

Values 0 to 720 seconds in 1 second increments

rapid-psc-timer

Syntax **rapid-psc-timer** *interval*
 no rapid-psc-timer

Context config>router>mpls>mpls-tp>protection-template

Description This command configures the rapid timer value to be used for protection switching coordination (PSC) packets for MPLS-TP linear protection (RFC 6378).

Default no rapid-psc-timer

Parameters *interval* — Specifies the rapid timer interval.

Values [10, 100, 1000 ms]

Default 10 ms

slow-psc-timer

Syntax **slow-psc-timer** *interval*
 no slow-psc-timer

Context config>router>mpls>mpls-tp>protection-template

Description This command configures the slow timer value to be used for protection switching coordination (PSC) packets for MPLS-TP linear protection (RFC 6378).

Default no rapid-psc-timer

Parameters	<i>interval</i> — Specifies the slow timer interval.
Values	[10, 100, 1000 ms]

global-id

Syntax	global-id <i>global-id</i> no global-id
Context	config>router>mpls>mpls-tp
Description	This command configures the MPLS-TP Global ID for the node. This is used as the 'from' Global ID used by MPLS-TP LSPs originating at this node. If a value is not entered, the Global ID is taken to be Zero. This is used if the global-id is not configured. If an operator expects that inter domain LSPs will be configured, then it is recommended that the global ID should be set to the local ASN of the node, as configured under config>system. If two-byte ASNs are used, then the most significant two bytes of the global-id are padded with zeros. In order to change the value of the global-id, config>router>mpls>mpls-tp must be in the shutdown state. This will bring down all of the MPLS-TP LSPs on the node. New values are propagated to the system when a no shutdown is performed.
Default	no global-id
Parameters	<i>global-id</i> — Specifies the global ID for the node.
Values	0 to 4294967295

node-id

Syntax	node-id <i>node-id</i> no node-id
Context	config>router>mpls>mpls-tp
Description	This command configures the MPLS-TP Node ID for the node. This is used as the 'from' Node ID used by MPLS-TP LSPs originating at this node. The default value of the node-id is the system interface IPv4 address. The Node ID may be entered in 4-octet IPv4 address format, <a.b.c.d>, or as an unsigned 32 bit integer. The Node ID is not treated as a routable IP address from the perspective of IP routing, and is not advertised in any IP routing protocols. The MPLS-TP context cannot be administratively enabled unless at least a system interface IPv4 address is configured because MPLS requires that this value is configured.
Default	no node-id

Parameters	<i>node-id</i> — Specifies the MPLS-TP node ID for the node.
Values	<a.b.c.d> or [1to 4294967295]
Default	System interface IPv4 address

transit-path

Syntax	transit-path <i>path-name</i> no transit-path
Context	config>router>mpls>mpls-tp
Description	This command enables the configuration or editing of an MPLS-TP transit path at an LSR.
Default	no transit-path
Parameters	<i>path-name</i> — Specifies the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

path-id

Syntax	path-id { lsp-num <i>lsp-num</i> working-path protect-path [src-global-id <i>src-global-id</i>] src-node-id <i>src-node-id</i> src-tunnel-num <i>src-tunnel-num</i> [dest-global-id <i>dest-global-id</i>] dest-node-id <i>dest-node-id</i> [dest-tunnel-num <i>dest-tunnel-num</i>]} no path-id
Context	config>router>mpls>mpls-tp>transit-path
Description	This command configures path ID for an MPLS-TP transit path at an LSR. The path ID is equivalent to the MPLS-TP LSP ID and is used to generate the maintenance entity group intermediate point (MIP) identifier for the LSP at the LSR. A path-id must be configured for on-demand OAM to verify an LSP at the LSR. The path-id must contain at least the following parameters: lsp-num, src-node-id, src-global-id, src-tunnel-num, dest-node-id. The path-id must be unique on a node. It is recommended that this is also configured to be a globally unique value. The no form of the command removes the path ID from the configuration.
Default	no path-id
Parameters	<i>lsp-num</i> — Specifies the LSP number. Values 1 to 65535, or working path, or protect-path. A working-path is equivalent to a lsp-num of 1, and a protect-path is an lsp-num of 2.

src-global-id — Specifies the source global ID.

Values 0 to 4294967295

src-node-id — Specifies the source node ID.

Values a.b.c.d or 1 to 4294967295

src-tunnel-num — Specifies the source tunnel number.

Values 1 to 61440

dest-global-id — Specifies the destination global ID. If the destination global ID is not entered, then it is set to the same value as the source global ID.

Values 0 to 4294967295

dest-node-id — Specifies the destination node ID.

Values a.b.c.d or 1 to 4294967295

dest-tunnel-num — Specifies the destination tunnel number. If the destination tunnel number is not entered, then it is set to the same value as the source tunnel number.

Values 1 to 61440

forward-path

Syntax	[no] forward-path
Context	config>router>mpls>mpls-tp>transit-path
Description	This command enables the forward path of an MPLS-TP transit path to be created or edited. The forward path must be created before the reverse path. The no form of this command removes the forward path. The forward path cannot be removed if a reverse exists.
Default	no forward-path

reverse-path

Syntax	[no] reverse-path
Context	config>router>mpls>mpls-tp>transit-path
Description	This command enables the reverse path of an MPLS-TP reverse path to be created or edited. The reverse path must be created after the forward path. The no form of this command removes the reverse path. The reverse path must be removed before the forward path.

Default no reverse-path

in-label

Syntax **in-label** *in-label* **out-label** *out-label* **out-link** *if-name* [**next-hop** *next-hop*]
no in-label

Context config>router>mpls>mpls-tp>transit-path>forward-path
 config>router>mpls>mpls-tp>transit-path>reverse-path

Description This command configures the label mapping associated with a forward path or reverse path of an MPLS-TP transit path to be configured.

The incoming label, outgoing label and outgoing interface must be configured, using the **in-label**, **out-label** and **out-link** parameters. If the out-link refers to a numbered IP interface, the user may optionally configure the **next-hop** parameter and the system will determine the interface to use to reach the configured next-hop, but will check that the user-entered value for the *out-link* corresponds to the link returned by the system. If they do not correspond, then the path will not come up.

Default no in-label

Parameters *in-label* — Specifies the in label.
 Values 32 to 16415
out-label — Specifies the out label.
 Values 32 to 16415
if-name — Specifies the name of the outgoing interface use for the path.
next-hop — Specifies the next-hop.
 Values a.b.c.d

shutdown

Syntax **[no] shutdown**

Context config>router>mpls>mpls-tp>transit-path

Description This command administratively enables or disables an MPLS-TP transit path.

Default no shutdown

LSP Commands

lsp

Syntax	[no] lsp <i>lsp-name</i> [bypass-only p2mp-lsp mpls-tp <i>src-tunnel-num</i> sr-te]
Context	config>router>mpls
Description	This command creates an LSP that is either signaled dynamically by the router, or a statically provisioned MPLS-TP LSP. 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 for signaled LSPs, or at least one working path for MPLS-TP LSPs. 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 shutdown before it can be deleted. The LSP must also be unbound from all SDPs before it can be deleted.
Default	none
Parameters	<i>lsp-name</i> — 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 router selects it. If no manual bypass tunnel is found, the router dynamically signals a bypass LSP in the default behavior. The CLI for this feature includes a knob that provides the user with the option to disable dynamic bypass creation on a per node basis. p2mp-lsp — Defines an LSP as a point-to-multipoint LSP. The following parameters can be used with a P2MP LSP: adaptive, adspec, cspf, exclude, fast-reroute, from, hop-limit, include, metric, retry-limit, retry-timer, resignal-timer. The following parameters cannot be used with a P2MP LSP: primary, secondary, to, dest-global-id, dest-tunnel-number, working-tp-path, protect-tp-path. This option is not supported on the 7450 ESS. mpls-tp <i>src-tunnel-num</i> — Defines an LSP as an MPLS-TP LSP. The <i>src-tunnel-num</i> is a mandatory create time parameter for mpls-tp LSPs, and has to be assigned by the user based on the configured range of tunnel IDs. The following parameters can only be used with an MPLS-TP LSP: to, dest-global-id, dest-tunnel-number, working-tp-path, protect-tp-path. Other parameters defined for the above LSP types cannot be used.

adaptive

Syntax	[no] adaptive
Context	config>router>mpls>lsp config>router>mpls>lsp>primary config>router>mpls>lsp>primary-p2mp-instance config>router>mpls>lsp>secondary config>router>mpls>lsp-template
Description	This command enables the make-before-break functionality for an LSP or LSP path. When enabled for the LSP, make-before-break will be performed for primary path and all the secondary paths of the LSP. The config>router>mpls>lsp>primary-p2mp-instance>adaptive command is not supported on the 7450 ESS.
Default	adaptive

adspec

Syntax	[no] adspec
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	When enabled, the ADSPEC object will be included in RSVP messages for this LSP. The ADSPEC object is used by the ingress LER to discover the minimum value of the MTU for links in the path of the LSP. By default, the ingress LER derives the LSP MTU from that of the outgoing interface of the LSP path. A bypass LSP always signals the ADSPEC object since it protects both primary paths which signal the ADSPEC object and primary paths which do not. This means that MTU of LSP at ingress LER may change to a different value from that derived from the outgoing interface even if the primary path has ADSPEC disabled.
Default	no adspec — No ADSPEC objects are included in RSVP messages.

auto-bandwidth

Syntax	[no] auto-bandwidth
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command enables (and the no form disables) automatic adjustments of LSP bandwidth.

Auto-bandwidth at the LSP level cannot be executed unless **adaptive** is configured in the **config>router>mpls>lsp** context.

Default no auto-bandwidth

adjust-down

Syntax **adjust-down** *percent* [**bw** *mbps*]
no adjust-down

Context config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command configures the minimum threshold for decreasing the bandwidth of an LSP based on active measurement of LSP bandwidth.

The **no** form of this command is equivalent to adjust-down 5.

Default no adjust-down

Parameters *percent* — Specifies the minimum difference between the current bandwidth reservation of the LSP and the (measured) maximum average data rate, expressed as a percentage of the current bandwidth, for decreasing the bandwidth of the LSP.

Values 1 to 100

Default 5

mbps — Specifies the minimum difference between the current bandwidth reservation of the LSP and the (measured) maximum average data rate, expressed as an absolute bandwidth (Mb/s), for decreasing the bandwidth of the LSP.

Values 0 to 100000

Default 0

adjust-up

Syntax **adjust-up** *percent* [**bw** *mbps*]
no adjust-up

Context config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command configures the minimum threshold for increasing the bandwidth of an LSP based on active measurement of LSP bandwidth.

The **no** form of this command is equivalent to adjust-up 5.

Default no adjust-up

Parameters	<i>percent</i> — Specifies the minimum difference between the current bandwidth reservation of the LSP and the (measured) maximum average data rate, expressed as a percentage of the current bandwidth, for increasing the bandwidth of the LSP. Values 1-100 Default 5
	<i>mbps</i> — Specifies the minimum difference between the current bandwidth reservation of the LSP and the (measured) maximum average data rate, expressed as an absolute bandwidth (Mb/s), for increasing the bandwidth of the LSP Values 0 to 100000 Default 0

egress-statistics

Syntax	[no] egress-statistics
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command configures statistics in the egress data path of an originating LSP at a head-end node. The user must execute the no shutdown for this command to effectively enable statistics. The same set of counters is updated for packets forwarded over any path of the LSP and over the lifetime of the LSP. In steady state, the 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. LSP statistics are not collected on a dynamic or a static bypass tunnel itself. 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. 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. A consequence of this is that when the user enables statistics collection on an RSVP LSP which is also used for tunneling LDP FECs with the LDP over RSVP feature, then statistics will be collected on the RSVP LSP only. There will be no statistics collected from an LDP FEC tunneled over this RSVP LSP regardless if the user enabled statistics collection on this FEC. When, the user disables statistics collection on the RSVP LSP, then statistics collection, if enabled, will be performed on a tunneled LDP FEC. 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

entropy-label

Syntax	entropy-label { force-disable enable inherit }
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command configures the use of entropy labels for an LSP. The entropy label and entropy label indicator (ELI) require 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 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 provides a means to override the default MPLS behavior on a per-LSP basis. For auto-LSPs, it can only be configured in LSP templates of type one-hope-p2p and mesh-p2p. Under the LSP context, when the value of entropy-label is set to enable, the ingress LER will take into consideration what is signaled from the egress node for ELC when marking the NHLFE as entropy-label-capable. Since the value of entropy-label is set to enable at the LSP level, the system will always mark it in the TTM as entropy-label-capable regardless of the signaled value, in order 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 ELC. When the value of entropy-label changes at either the MPLS level or the LSP level, the new operational value will not take effect until the LSP is resigaled. A shutdown/no shutdown of the LSP is required to enable the new value. The user can use the clear command or bounce MPLS itself (shutdown/no shutdown) to force the new value to take effect for a large numbers of LSPs.
Default	entropy-label inherit
Parameters	force-disable — Indicates that 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.

- enable** — Indicates that 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. Therefore, 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.
- inherit** — Indicates that the value of **entropy-label** is inherited from the setting in the MPLS context.

fc

- Syntax** **fc** *fc-name* **sampling-weight** *sampling-weight*
 no fc
- Context** config>router>mpls>lsp-template>auto-bandwidth
- Description** This command configures the sampling weight.

max-bandwidth

- Syntax** **max-bandwidth** *mbps*
 no max-bandwidth
- Context** config>router>mpls>lsp>auto-bandwidth
 config>router>mpls>lsp-template>auto-bandwidth
- Description** This command configures the maximum bandwidth that auto-bandwidth allocation is allowed to request for an LSP.
- The LSP maximum applies whether the bandwidth adjustment is triggered by normal adjust-interval expiry, the overflow limit having been reached, or manual request.
- The **no** form of the command means max-bandwidth is 100 Gbps.
- The max-bandwidth must be greater than the min-bandwidth.
- Default** no max-bandwidth
- Parameters** *mbps* — Specifies the maximum bandwidth in Mb/s.
- Values** 0 to 100000
- Default** 0

min-bandwidth

- Syntax** **min-bandwidth** *mbps*
 no min-bandwidth

Context	config>router>mpls>lsp>auto-bandwidth config>router>mpls>lsp-template>auto-bandwidth				
Description	This command configures the minimum bandwidth that auto-bandwidth allocation is allowed to request for an LSP. The LSP minimum applies whether the bandwidth adjustment is triggered by normal adjust-timer expiry or manual request. The no form of the command means min-bandwidth is zero.				
Default	no min-bandwidth				
Parameters	<i>mbps</i> — Specifies the minimum bandwidth in Mb/s. <table> <tr> <td>Values</td><td>0 to 100000</td></tr> <tr> <td>Default</td><td>0</td></tr> </table>	Values	0 to 100000	Default	0
Values	0 to 100000				
Default	0				

monitor-bandwidth

Syntax	[no] monitor-bandwidth
Context	config>router>mpls>lsp>auto-bandwidth config>router>mpls>lsp-template>auto-bandwidth
Description	This command enables the collection and display of auto-bandwidth measurements, but prevents any automatic bandwidth adjustments from taking place. This command is mutually exclusive with the overflow-limit command. The no form of the command the collection and display of auto-bandwidth measurements.

multipliers

Syntax	multipliers sample-multiplier <i>num1</i> adjust-multiplier <i>num2</i> no multipliers
Context	config>router>mpls>lsp>auto-bandwidth config>router>mpls>lsp-template>auto-bandwidth
Description	This command configures the sample-multiplier and adjust-multiplier applicable to one particular LSP.

The sample-multiplier configures the number of collection intervals between measurements of the number of bytes that have been transmitted on the LSP. The byte counts include the layer 2 encapsulation of MPLS packets and represent traffic of all forwarding classes and priorities (in-profile vs, out-of-profile) belonging to the LSP. The router calculates the average data rate in each sample interval. The maximum of this average data rate over multiple sample intervals is the measured bandwidth input to the auto-bandwidth adjustment algorithms.

The adjust-multiplier is the number of collection intervals between periodic evaluations by the ingress LER about whether to adjust the LSP bandwidth. The router keeps track of the maximum average data rate of each LSP since the last reset of the adjust-count.

The adjust-multiplier is not allowed to be set to a value less than the sample-multiplier. It is recommended that the adjust-multiplier be a multiple of the sample-multiplier.

The **no** form of this command instructs the system to take the value from the auto-bandwidth-defaults command.

Default no multipliers

Parameters *number1* — The number of collection intervals in a sample interval.

Values 1 to 511

Default inherited

number2 — The number of collection intervals in an adjust interval.

Values 1 to 16383

Default inherited

overflow-limit

Syntax **overflow-limit** *number threshold percent [bw mbps]*
no overflow-limit

Context config>router>mpls>lsp>auto-bandwidth
config>router>mpls>lsp-template>auto-bandwidth

Description This command configures overflow-triggered auto-bandwidth adjustment. It sets the threshold at which bandwidth adjustment is initiated due to the configured number of overflow samples having been reached, regardless of how much time remains until the adjust interval ends.

A sample interval is counted as an overflow if the average data rate during the sample interval is higher than the currently reserved bandwidth by at least the thresholds configured as part of this command.

If overflow-triggered auto-bandwidth adjustment is successful the overflow count, maximum average data rate and adjust count are reset. If overflow-triggered auto-bandwidth adjustment fails then the overflow count is reset but the maximum average data rate and adjust count maintain current values.

This command is mutually exclusive with the `monitor-bandwidth` command.

The **no** form of this command disables overflow-triggered automatic bandwidth adjustment.

Default	no overflow-limit
Parameters	<i>number</i> — The number of overflow samples that triggers an overflow auto-bandwidth adjustment attempt. Values 1 to 10 Default 0 (disabled) <i>percent</i> — The minimum difference between the current bandwidth of the LSP and the sampled data rate, expressed as a percentage of the current bandwidth, for counting an overflow sample. Values 1 to 100 Default 0 (disabled) <i>mbps</i> — The minimum difference between the current bandwidth of the LSP and the sampled data rate, expressed as an absolute bandwidth (Mb/s) relative to the current bandwidth, for counting an overflow sample. Values 1to 100000 Default 0 (disabled)

underflow-limit

Syntax	underflow-limit <i>number threshold percent [bw mbps]</i> no underflow-limit
Context	config>router>mpls>lsp>auto-bandwidth config>router>mpls>lsp-template>auto-bandwidth
Description	This command configures underflow-triggered auto-bandwidth adjustment. An underflow auto-bandwidth adjustment can occur any time during the adjust-interval; it is triggered when the number of consecutive underflow samples reaches the threshold N configured as part of this command. The new bandwidth of the LSP after a successful underflow adjustment is the maximum data rate observed in the last N consecutive underflow samples. A sample interval is counted as an underflow if the average data rate during the sample interval is lower than the currently reserved bandwidth by at least the thresholds configured as part of this command. This command is mutually exclusive with the monitor-bandwidth command.

The **no** form of this command disables underflow-triggered automatic bandwidth adjustment.

Default	no underflow-limit
Parameters	<i>number</i> — The number of consecutive underflow samples that triggers an underflow auto-bandwidth adjustment attempt. Values 0 to 10 Default 0 (disabled) <i>percent</i> — The minimum difference between the current bandwidth of the LSP and the sampled data rate, expressed as a percentage of the current bandwidth, for counting an underflow sample. Values 0 to 100 Default 0 (disabled) <i>mbps</i> — The minimum difference between the current bandwidth of the LSP and the sampled data rate, expressed as an absolute bandwidth (Mb/s) relative to the current bandwidth, for counting an underflow sample. Values 0 to 1000000 Default 0 (disabled)

bfd

Syntax	bfd
Context	config>router>mpls>lsp config>router>mpls>lsp>primary config>router>mpls>lsp-template
Description	The bfd command creates a context for the configuration of LSP BFD commands.

bfd-enable

Syntax	bfd-enable no bfd-enable
Context	config>router>mpls>lsp>bfd config>router>mpls>lsp>primary>bfd config>router>mpls>lsp-template>bfd
Description	This command enables LSP BFD on the LSP. Lsp-bfd must also be configured under config>router to enable LSP BFD. The parameters for the BFD session are derived from the named BFD Template, which must have been configured prior to the bfd-enable command and associated with the service using the bfd-template command.
Default	no bfd-enable

bfd-template

Syntax	bfd-template <i>name</i> no bfd-template
Context	config>router>mpls>lsp>bfd config>router>mpls>lsp>primary>bfd config>router>mpls>lsp-template>bfd
Description	This command references a named BFD template to be used by LSP BFD. The template specifies parameters, such as the minimum transmit and receive control packet timer intervals, to be used by the BFD session. Templates are configured under the config>router>bfd context.
Default	no bfd-template
Parameters	<i>name</i> — Specifies a text string name for the template of up to 32 characters in printable 7-bit ASCII, enclosed in double quotes. Default none

lsp-ping-interval

Syntax	lsp-ping-interval <i>seconds</i> no lsp-ping-interval
Context	config>router>mpls>lsp>bfd config>router>mpls>lsp>primary>bfd config>router>mpls>lsp-template>bfd
Description	This command configures the interval for the periodic LSP ping for LSPs on which bfd-enable has been configured. This is used to bootstrap and maintain the LSP BFD session. The default interval is 60 seconds, with a maximum of 300 seconds. A value of 0 disables periodic LSP Ping, such that an LSP Ping containing a bootstrap TLV is only sent when the BFD session is first initialized. In scaled environments, LSP BFD sessions should use longer timers to reduce the chance of congestion and loading of common resources. Unless required, the lsp-ping-interval should not be set lower than 300 seconds.
Default	no lsp-ping interval This sets the periodic LSP Ping interval to a default of 60 s.
Parameters	<i>seconds</i> — Sets the periodic LSP Ping interval. Values 0, 60 to 300 s Default 60 s

bgp-shortcut

Syntax	[no] bgp-shortcut
Context	config>router>mpls>lsp
Description	This command enables the use of RSVP LSP for IPv4 BGP routes.

bgp-transport-tunnel

Syntax	bgp-transport-tunnel [include exclude]
Context	config>router>mpls>lsp
Description	This command allows or blocks RSVP-TE LSP to be used as a transport LSP for BGP tunnel routes.
Default	bgp-transport-tunnel include
Parameters	include — Allows RSVP-TE LSP to be used as transport LSP from the ASBR to local PE router, from ingress PE to ASBR in the local AS or between multi-hop eBGP peers with ASBR to ASBR adjacency. exclude — Blocks RSVP-TE LSP to be used as transport LSP from the ASBR to local PE router, from ingress PE to ASBR in the local AS or between multi-hop eBGP peers with ASBR to ASBR adjacency.

class-forwarding

Syntax	[no] class-forwarding
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command enables the context to configure class based forwarding parameters for a given LSP or LSP-template. A change in the Class-Based Forwarding configuration may result in a change of forwarding behavior. The no form removes any Class-Based Forwarding configuration associated to that LSP or LSP-template.
Default	no class-forwarding

fc

Syntax	fc { be l2 af l1 h2 ef h1 nc } no fc [{ be l2 af l1 h2 ef h1 nc }]
Context	config>router>mpls>lsp>class-forwarding config>router>mpls>lsp-template>class-forwarding
Description	This command assigns a forwarding class to a given LSP or LSP-template. This command can only be passed with a single forwarding class but by passing the command multiple times it is possible to assign multiple forwarding classes (up to 8) to the same LSP or LSP-template. A change in the Class-Based Forwarding configuration may result in a change of forwarding behavior. The no form of this command removes the assignment of the forwarding classes from the LSP or LSP-template. It can only be passed with either a single or no forwarding class. If no forwarding class is specified, all the assignments are removed. In the other case, only the assignment of the specified forwarding class is removed.
Default	no fc

default-lsp

Syntax	[no] default-lsp
Context	config>router>mpls>lsp>class-forwarding config>router>mpls>lsp-template>class-forwarding
Description	This command assigns the default-lsp configuration to a given LSP or LSP-template. The Default LSP is the LSP on which will be forwarded any packet associated to a given class but for which no LSP with the corresponding class explicitly assigned exists. A change in the Class-Based Forwarding configuration may result in a change of forwarding behavior. The no form of this command removes the default-lsp assignment from the LSP or LSP-template.
Default	no default-lsp

class-type

Syntax	class-type <i>ct-number</i> no class-type
Context	config>router>mpls>lsp config>router>mpls>lsp>primary

config>router>mpls>lsp>secondary

Description	This command configures the Diff-Serv Class Type (CT) for an LSP, the LSP primary path, or the LSP secondary path. The path level configuration overrides the LSP level configuration. However, only one CT per LSP path will be allowed as per RFC 4124. The signaled CT of a dynamic bypass is always be CT0 regardless of the CT of the primary LSP path. The setup and hold priorities must be set to default values, i.e., 7 and 0 respectively. This assumes that the operator configured a couple of TE classes, one which combines CT0 and a priority of 7 and the other which combines CT0 and a priority of 0. If not, the bypass LSP will not be signaled and will go into the down state. The operator cannot configure the CT, setup priority, and hold priority of a manual bypass. They are always signaled with CT0 and the default setup and holding priorities. The signaled CT and setup priority of a detour LSP must match those of the primary LSP path it is associated with. If the operator changes the CT of an LSP or of an LSP path, or changes the setup and holding priorities of an LSP path, the path will be torn down and retried. An LSP which does not have the CT explicitly configured will behave like a CT0 LSP when Diff-Serv is enabled. If the operator configured a combination of a CT and a setup priority and/or a combination of a CT and a holding priroty for an LSP path that are not supported by the user-defined TE classes, the LSP path will be kept in a down state and an error code will be displayed in the show command output for the LSP path. The no form of this command reverts to the default value.				
Default	no class-type				
Parameters	<i>ct-number</i> — The Diff-Serv Class Type number. <table><tr><td>Values</td><td>0 – 7</td></tr><tr><td>Default</td><td>0</td></tr></table>	Values	0 – 7	Default	0
Values	0 – 7				
Default	0				

bandwidth

Syntax	bandwidth <i>rate-in-mbps</i>
Context	config>router>mpls>lsp>primary-p2mp-instance config>router>mpls>lsp>template
Description	This command specifies the amount of bandwidth to be reserved for the P2MP instance. The config>router>mpls>lsp>primary-p2mp-instance>bandwidth command is not supported on the 7450 ESS.

Parameters *rate-in-mbps* — specifies the bandwidth, in Mb/s.

Values 0 to 100000

cspf

Syntax **[no] cspf [use-te-metric]**

Context config>router>mpls>lsp
config>router>mpls>lsp-template

Description This command enables Constrained Shortest Path First (CSPF) computation for constrained-path LSPs. Constrained-path LSPs are the ones that take configuration constraints into account. CSPF is also used to calculate the detour routes when fast-reroute is enabled.

Explicitly configured LSPs where each hop from ingress to egress is specified do not use CSPF. The LSP will be set up using RSVP signaling from ingress to egress.

If an LSP is configured with **fast-reroute** *frr-method* specified but does not enable CSPF, then neither global revertive nor local revertive will be available for the LSP to recover.

Default no cspf

Parameters **use-te-metric** — Specifies to use the use of the TE metric for the purpose of the LSP path computation by CSPF.

dest-global-id

Syntax **dest-global-id** *dest-global-id*
no dest-global-id

Context config>router>mpls>lsp

Description This optional command configures the MPLS-TP Global ID of the far end node of the MPLS-TP LSP. This command is only allowed for MPLS-TP LSPs. Global ID values of 0 indicate that the local node's configured global ID is used. If the local global-id is 0, then the dest-global-id must also be 0. The dest-global-id cannot be changed if an LSP is in use by an SDP.

Default 0

Parameters *dest-global-id* — Specifies the destination global ID.

Values 0 to 4294967295

Default 0

dest-tunnel-number

Syntax	dest-tunnel-number <i>dest-tunnel-number</i> no dest-tunnel-number
Context	config>router>mpls>lsp
Description	This optional command configures the MPLS-TP tunnel number of the LSP at the far end node of the MPLS-TP LSP. This command is only allowed for MPLS-TP LSPs. If it is not entered, then the system will take the dest-tunnel-number to be the same as the src-tunnel-num for the LSP.
Default	The default value is the configured <i>src-tunnel-num</i> .
Parameters	<i>dest-tunnel-number</i> — Specifies the destination tunnel number.
Values	1 to 61440
Default	<i>src-tunnel-number</i>

working-tp-path

Syntax	[no] working-tp-path
Context	config>router>mpls>lsp
Description	This command creates or edits the working path for an MPLS-TP LSP. At least one working path (but not more than one working path) must be created for an MPLS-TP LSP. If MPLS-TP linear protection is also configured, then this is the path that is used as the default working path for the LSP, and it must be created prior to the protect path. The working-tp-path can only be deleted if no protect-tp-path exists for the LSP. The following commands are applicable to the working-tp-path: lsp-num, in-label, out-label, mep, shutdown.
Default	no working-tp-path

protect-tp-path

Syntax	[no] protect-tp-path
Context	config>router>mpls>lsp
Description	This command creates or edits the protect path for an MPLS-TP LSP. At least one working path must exist before a protect path can be created for an MPLS-TP LSP. If MPLS-TP linear protection is also configured, then this is the path that is used as the default protect path for the LSP. The protect path must be deleted before the wokring path. Only one protect path can be created for each MPLS-TP LSP.

The following commands are applicable to the working-tp-path: **lsp-num**, **in-label**, **out-label**, **mep**, **shutdown**.

lsp-num

Syntax	lsp-num <i>lsp-num</i> no lsp-num				
Context	config>mpls>lsp>working-tp-path config>mpls>lsp>protect-tp-path				
Description	This command configures the MPLS-TP LSP Number for the working TP path or the Protect TP Path.				
Default	no lsp-num				
Parameters	<i>lsp-num</i> — Specifies the LSP number. <table> <tr> <td>Values</td><td>1 to 65535</td></tr> <tr> <td>Default</td><td>1 for a working path, 2 for a protect path</td></tr> </table>	Values	1 to 65535	Default	1 for a working path, 2 for a protect path
Values	1 to 65535				
Default	1 for a working path, 2 for a protect path				

in-label

Syntax	in-label <i>in-label</i> no in-label		
Context	config>mpls>lsp>working-tp-path config>mpls>lsp>protect-tp-path		
Description	This command configures the incoming label for the reverse path or the working path or the protect path of an MPLS-TP LSP. MPLS-TP LSPs are bidirectional, and so an incoming label value must be specified for each path.		
Default	no in-label		
Parameters	<i>in-label</i> — Specifies the in label. <table> <tr> <td>Values</td><td>32 to 16415</td></tr> </table>	Values	32 to 16415
Values	32 to 16415		

out-label

Syntax	out-label <i>out-label</i> out-link <i>if-name</i> [next-hop <i>ip-address</i>] no out-label
Context	config>mpls>lsp>working-tp-path config>mpls>lsp>protect-tp-path

Description	This command configures the outgoing label value to use for an MPLS-TP working or protect path. The out-link is the outgoing interface on the node that this path will use, and must be specified. If the out-link refers to a numbered IP interface, the user may optionally configure the next-hop parameter and the system will determine the interface to use to reach the configured next-hop, but will check that the user-entered value for the <i>out-link</i> corresponds to the link returned by the system. If they do not correspond, then the path will not come up.
Default	no out-label
Parameters	<i>out-label</i> — Specifies the out label. Values 32 to 16415 <i>if-name</i> — Specifies the interface name. <i>ip-address</i> — Specifies the IPv4 address in a.b.c.d

mep

Syntax	[no] mep
Context	config>mpls>lsp>working-tp-path config>mpls>lsp>protect-tp-path
Description	This command creates or edits an MPLS-TP maintenance entity group (MEG) endpoint (MEP) on an MPLS-TP path. MEPs represent the termination point for OAM flowing on the path, as well as linear protection for the LSP. Only one MEP can be configured at each end of the path. The following commands are applicable to a MEP on an MPLS-TP working or protect path: oam-template, bfd-enable, and shutdown. In addition, a protection-template may be configured on a protect path. The no form of the command removes a MEP from an MPLS-TP path.

mip

Syntax	[no] mip
Context	config>router>mpls>lsp>transit-path>forward-path config>router>mpls>lsp>transit-path>reverse-path
Description	This command creates a context for maintenance entity group intermediate point (MIP) parameters for the forward path and the reverse path of an MPLS-TP LSP at an LSR.
Default	none

bfd-trap-suppression

Syntax	[no] bfd-trap-suppression
Context	config>router>mpls>lsp>protect-tp-path>mep config>router>mpls>lsp>working-tp-path>mep
Description	This command enables AIS packets on a working or protection path of an MPLS-TP LSP to suppress BFD Down traps if a BFD session goes down on that path. It also causes BFD Up traps to be suppressed, and enables the 2.5 s hold-down timer. Suppression only occurs as a result of a received AIS packet. Traps generated as a result of a local failure at an LER are not suppressed. The no form of the command disables BFD Down/Up trap suppression when AIS packets are received.
Default	no bfd-trap-suppression

dsmap

Syntax	dsmap <i>if-num</i> no dsmap
Context	config>router>mpls>lsp>working-tp-path>mep config>router>mpls>lsp>protect-tp-path>mep config>router>mpls>lsp>transit-path>forward-path>mip config>router>mpls>lsp>transit-path>reverse-path>mip
Description	This command is used to configure the values to use in the DSMAP TLV sent by a node in an LSP Trace echo request for a static MPLS-TP LSP. A node sending a DSMAP TLV will include the in-if-num and out-if-num values. Additionally, it will include the out-label for the LSP in the Label TLV for the DSMAP in the echo request message.
Parameters	<i>if-num</i> — This is a 32-bit value corresponding to the expected ingress interface if-num used by an MPLS-TP LSP for the next hop downstream. A value of zero means that no interface validation will be performed. Values 0 to 4294967295 Default 0

oam-template

Syntax	oam-template <i>name</i> no oam-template
Context	config>mpls>lsp>working-tp-path config>mpls>lsp>protect-tp-path

Description	This command applies a OAM template to an MPLS-TP working or protect path. It contains configuration parameters for proactive OAM mechanisms that can be enabled on the path e.g. BFD. Configuration of an OAM template is optional. The no form of the command removes the OAM template from the path.
Default	no oam-template
Parameters	<i>name</i> — Specifies a text string name for the template up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

bfd-enable

Syntax	bfd-enable [cc cc_cv] no bfd-enable
Context	config>mpls>lsp>working-tp-path config>mpls>lsp>protect-tp-path
Description	The command associates the operational state of an MPLS-TP path with a BFD session whose control packets flow on the path. The BFD packets are encapsulated in a generic associated channel (G-ACh) on the path. The timer parameters of the BFD session are taken from the the OAM template of the MEP. A value of cc means that the BFD session is only used for continuity check of the the MPLS-TP path. In this case, the cc timer parameters of the OAM template apply. A value of cv means that the BFD session is used for both continuity checking and connectivity verification, and the cv timers of the OAM template apply. This form of the bfd-enable command is only applicable when it is configured under a MEP used on an MPLS-TP working or protect path.
Default	no bfd-enable
Parameters	cc cc_cv — cc indicates that BFD runs in CC only mode. This mode uses GACH channel type 0x07. cc_cv indicates that BFD runs in combined CC and CV mode. This mode uses channel type 0x22 for MPLS-TP CC packets, and 0x23 for MPLS-TP CV packets.

protection-template

Syntax	protection-template <i>name</i> no protection-template
Context	config>mpls>lsp>protect-tp-path

Description	This command applies a protection template name to an MPLS-TP LSP that the protect path is configured under. If the template is applied, then MPLS-TP 1:1 linear protection is enabled on the LSP, using the parameters specified in the named template. A named protection template can only be applied to the protect path context of an MPLS-TP LSP. The no form of the command removes the template and thus disables mpls-tp linear protection on the LSP.
Default	no protection-template
Parameters	<i>name</i> — Specifies at text string for the template up to 32 characters in printable 7-bit ASCII, enclosed in double quotes.

exclude

Syntax	[no] exclude <i>group-name</i> [<i>group-name...</i> (up to 5 max)]
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command specifies the admin groups to be excluded when an LSP is set up in the primary or secondary contexts. Each single operation of the exclude command allows a maximum of 5 groups to be specified at a time. However, a maximum of 32 groups can be specified per LSP through multiple operations. The admin groups are defined in the config>router>if-attribute>admin-group context. The exclude statement instructs the CSPF algorithm to avoid TE links which belong to any of the specified admin groups. A link which belongs one or more of the specified admin groups is excluded and thus pruned from the TE database before the CSPF computation. Use the no form of the command to remove the exclude command.
Default	no exclude
Parameters	<i>group-name</i> — Specify the existing group-name to be excluded when an LSP is set up.

exclude-node

Syntax	[no] exclude-node <i>ip-address</i>
Context	config>router>mpls>lsp
Description	This command enables the option to include XRO object in the bypass LSP PATH message object. The exclude-node option is required for manual bypass LSP with XRO to FRR protect ABR node in a multi-vendor network depolyment. This command must be configured on the PLR node that protects the ABR node. The ABR node IP address must be configured as exclude-node.

Default no exclude-node

fast-reroute

Syntax **fast-reroute** *frr-method*
 no fast-reroute

Context config>router>mpls>lsp
 config>router>mpls>lsp-template

Description This command creates a pre-computed detour LSP from each node in the path of the LSP. In case of failure of a link or LSP between two nodes, traffic is immediately rerouted on the pre-computed detour LSP, thus avoiding packet-loss.

When **fast-reroute** is enabled, each node along the path of the LSP tries to establish a detour LSP as follows:

- Each upstream node sets up a detour LSP that avoids only the immediate downstream node, and merges back on to the actual path of the LSP as soon as possible.
If it is not possible to set up a detour LSP that avoids the immediate downstream node, a detour can be set up to the downstream node on a different interface.
- The detour LSP may take one or more hops (see [hop-limit](#)) before merging back on to the main LSP path.
- When the upstream node detects a downstream link or node failure, the ingress router switches traffic to a standby path if one was set up for the LSP.

Fast reroute is available only for the primary path. No configuration is required on the transit hops of the LSP. The ingress router will signal all intermediate routers using RSVP to set up their detours. TE must be enabled for fast-reroute to work.

If an LSP is configured with **fast-reroute** *frr-method* specified but does not enable CSPF, then neither global revertive nor local revertive will be available for the LSP to recover.

The **no** form of the **fast-reroute** command removes the detour LSP from each node on the primary path. This command will also remove configuration information about the hop-limit and the bandwidth for the detour routes.

The **no** form of **fast-reroute hop-limit** command reverts to the default value.

Default no fast-reroute — When fast-reroute is specified, the default fast-reroute method is one-to-one.

Parameters	frr-method — Configures the fast-reroute method.
Values	one-to-one — In the one-to-one technique, a label switched path is established which intersects the original LSP somewhere downstream of the point of link or node failure. For each LSP which is backed up, a separate backup LSP is established.
Values	facility — This option, 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 which serves to backup up a set of LSPs. This LSP tunnel is called a bypass tunnel.
Values	The bypass tunnel must intersect the path of the original LSP(s) somewhere downstream of the point of local repair (PLR). Naturally, this constrains the set of LSPs being backed-up via that bypass tunnel to those that pass through a common downstream node. All LSPs which pass through the PLR and through this common node which do not also use the facilities involved in the bypass tunnel are candidates for this set of LSPs.

hop-limit

Syntax	hop-limit <i>limit</i> no hop-limit
Context	config>router>mpls>lsp>fast-reroute config>router>mpls>lsp-template>fast-reroute
Description	For fast reroute, how many more routers a detour is allowed to traverse compared to the LSP itself. For example, if an LSP traverses four routers, any detour for the LSP can be no more than ten router hops, including the ingress and egress routers.
Default	16
Parameters	<i>limit</i> — Specify 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.

Default	node-protect

from

Syntax	from <i>ip-address</i>

Context	config>router>mpls>lsp config>router>mpls>lsp-template

Description	This optional 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.

If an interface IP address is specified as the **from** address, and the egress interface of the nexthop 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	The system IP address

Parameters	<i>ip-address</i> — This is 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.

Default	System IP address

Values	System IP or network interface IP addresses

hop-limit

Syntax	hop-limit <i>number</i> no hop-limit

Context	config>router>mpls>lsp config>router>mpls>lsp>primary-p2mp-instance

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 with the following implications.

If the new value is less than the current number of hops of the established LSP, the LSP is brought down. The software then tries to re-establish the LSP within the new **hop-limit** number. If the new value is equal to or greater than the current number hops of the established LSP, the LSP is not affected.

The `config>router>mpls>lsp>primary-p2mp-instance>hop-limit` command is not supported on the 7450 ESS.

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

Default	255
Parameters	<i>number</i> — 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 [<i>offset</i>]] [no] igp-shortcut
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command enables the use of a specific RSVP LSP by IS-IS and OSPF routing protocols as a shortcut or as a forwarding adjacency for resolving IGP routes.

When the **rsvp-shortcut** or the advertise-tunnel-link option is enabled at the IGP instance level, all RSVP LSPs originating on this node are eligible by default as long as the destination address of the LSP, as configured in **config>router>mpls>lsp>to**, corresponds to a router-id of a remote node.

The **lfa-protect** option allows an LSP to be included in both the main SPF and the Loop-Free Alternate (LFA) SPF. 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 computation selected a tunneled primary next-hop for a prefix, the LFA SPF will not select an LFA next-hop for this prefix and the protection of this prefix will rely on the RSVP LSP FRR protection. If the main SPF computation selected a direct primary next-hop, then the LFA SPF will select an LFA next-hop for this prefix but will prefer a direct LFA next-hop over a tunneled LFA next-hop.

The **lfa-only** option allows an LSP to be included in the LFA SPF only such that the introduction of IGP shortcuts does not impact the main SPF decision. For a given prefix, the main SPF always selects a direct primary next-hop. The LFA SPF will select a an LFA next-hop for this prefix but will prefer a direct LFA next-hop over a tunneled LFA next-hop.

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 computing the cost of a prefix which is resolved to the LSP. The offset value is optional and it defaults to zero. The minimum net cost for a prefix is one (1) after applying the offset. The TTM continues to show the LSP operational metric as provided by MPLS. In other words, applications such as LDP-over-RSVP (when IGP shortcut is disabled) and BGP and static route shortcuts will continue to use the LSP operational metric.

The **relative-metric** option is mutually exclusive with the **lfa-protect** or the **lfa-only** options. In other words, an LSP with the **relative-metric** option enabled cannot be included in the LFA SPF and vice-versa when the **rsvp-shortcut** option is enabled in the IGP.

Finally, the **relative-metric** option is ignored when forwarding adjacency is enabled in IS-IS or OSPF. 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 maximum link metric allowed in that IGP. Both the main SPF and the LFA SPFs will use the local IGP database to resolve the routes.

The **no** form of this command disables the use of a specific RSVP LSP by IS-IS and OSPF routing protocols as a shortcut or a forwarding adjacency for resolving IGP routes.

Default	igp-shortcut. All RSVP 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. lfa-only — An LSP is included in the LFA SPF only. relative-metric [<i>offset</i>] — 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 an integer and is optional. An offset value of zero is used when the relative-metric option is enabled without specifying the offset parameter value. Values [-10, +10]

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 will find all equal cost shortest paths which satisfy the constraints of this path. Then, CSPF identifies the single link in each of these paths which has the least available bandwidth as a percentage of its maximum reservable bandwidth. It then selects the path which has the largest value of this percentage least available bandwidth figure. 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.

CSPF applies the least-fill criterion separately to each pre-emption priority in the base TE. A higher setup priority path can pre-empt lower holding priority paths.

CSPF also applies the least-fill criterion separately to each Diff-Serv TE class if Diff-Serv TE is enabled on this node. A higher setup priority path can pre-empt lower holding priority paths within a Class Type.

MPLS will re-signal and move the LSP to the new path in the following cases:

- Initial LSP path signaling.
- Re-try of an LSP path after failure.
- Make-before-break (MBB) due to pending soft preemption of the LSP path.
- MBB due to LSP path configuration change, i.e., a user change to bandwidth parameter of primary or secondary path, or a user enabling of fast-reroute option for the LSP.
- MBB of secondary path due to an update to primary path SRLG.
- MBB due to FRR Global Revertive procedures on the primary path.
- Manual re-signaling of an LSP path or of all LSP paths by the user.

During a manual re-signaling of an LSP path, MPLS will always re-signal the path regardless of whether the new path is exactly the same or different than the current path and regardless of whether the metric of the new path is different or not from that of the current path.

During a timer-based re-signaling of an LSP path which has the least-fill option enabled, MPLS will only re-signal the path if the metric of the new path is different than the one of the current path.

The user deletes a specific node entry in this database by executing the no form of this command.

Default no least-fill. The path of an LSP is randomly chosen among a set of equal cost paths.

ingress-statistics

Syntax [no] ingress-statistics

Context config>router>mpls>lsp

Description	This command configures the LSP ingress statistics.

ldp-over-rsvp

Syntax	[no] ldp-over-rsvp [include exclude]
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command configures an LSP so that it can be used by the IGP to calculate its SPF tree. When the ldp-over-rsvp option is also enabled in ISIS or OSPF, the IGP provides LDP with all ECMP IP next-hops and tunnel endpoints that it considers to be the lowest cost path to its destination. IGP provides only the endpoints which are the closest to the destination in terms of IGP cost for each IP next-hop of a prefix. If this results in more endpoints than the ECMP value configured on the router, it will further prune the endpoints based on the lowest router-id and for the same router-id, it will select lowest interface-index first. LDP then looks up the tunnel table to select the actual tunnels to the endpoint provided by IGP and further limits the endpoint selection to the ones which are the closest to destination across all the IP next-hops provided by IGP for a prefix. For each remaining endpoint, LDP selects a tunnel in a round-robin fashion until the router ECMP value is reached. For each endpoint, only tunnels with the same lowest metric are candidates. If more than one tunnel qualifies, the selection begins with the lowest tunnel-id. The no form of the command reverts to default operation.
Default	ldporsvp include

include

Syntax	[no] include <i>group-name</i> [<i>group-name...</i> (up to 5max)]
Context	config>router>mpls>lsp config>router>mpls>lsp>primary config>router>mpls>lsp>primary-p2mp-instance 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, up to 32 maximum. The include statement instructs the CSPF algorithm to pick TE links among the links which belong to one or more of the specified admin groups. A link that does not belong to at least one of the specified admin groups is excluded and thus pruned from the TE database before the CSPF computation. However, a link can still be selected if it belongs to one of the groups in a include statement but also belongs to other groups which are not part of any include statement in the LSP or primary/secondary path configuration. In other words, the include statements implements the "include-any" behavior. The config>router>mpls>lsp>primary-p2mp-instance>include command is not supported on the 7450 ESS. The no form of the command deletes the specified groups in the specified context.
Default	no include
Parameters	<i>group-name</i> — Specifies admin groups to be included when an LSP is set up.

path-profile

Syntax	path-profile <i>profile-id</i> [path-group <i>group-id</i>] no path-profile <i>profile-id</i>
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 and/or terminating on the same or different PE routers. In order to indicate this constraint to PCE, the user must configure the PCE path profile ID and Path Group ID the PCE computed or PCE controlled LSP belongs to. These parameters are passed transparently by PCC to PCE and are thus opaque data to the router. The association of the optional path-group ID is to allow PCE 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 5 entries of path-profile [path-group] can be associated with the same LSP.
Parameters	<i>profile-id</i> — Specifies the profile ID. Values 1 to 4294967295 path-group <i>group-id</i> — integer. Values 0 to 4294967295

pce-computation

Syntax	[no] pce-computation
Context	config>router>mpls>lsp
Description	This command enables a PCE computed LSP mode of operation. The user can grant the path computation requests only (PCE-computed) or both path computation requests and path update (PCE-controlled) to a PCE for a specific LSP. The pce-computation option indicates the path computation request is sent to the PCE instead of the local CSPF. Enabling it means the PCE can perform path computations for the LSP at the request of the router only. This is used in cases where the operator wants to make use of the PCE specific path computation algorithm instead of the local router CSPF algorithm. The default value is no pce-computation. Enabling this command is conditional on the cspf option being also enabled, otherwise the command is rejected. If the user disabled cspf option for an LSP the pce-computation option will also be automatically disabled. Enabling cspf option without enabling pce-computation for a SR-TE LSP means that internally the router is still doing label translation as if cspf was disabled since there is no support of CSPF for a SR-TE LSP on the router.
Default	no pce-computation

pce-control

Syntax	[no] pce-control
Context	config>router>mpls>lsp
Description	This command enables a PCE controlled LSP mode of operation. The pce-control option means the router delegates full control of the LSP to the PCE (PCE controlled). Enabling it means the PCE is acting in stateful-active mode for this LSP and PCE will be able to reroute the path following a failure or to re-optimize the path and update the router without the router requesting it. Note the user can delegate CSPF and non-CSPF LSPs. Also, the user can delegate LSPs which have the pce-computation option enabled or disabled. The LSP maintains its latest active path computed by PCE or the router at the time it is delegated. The PCE will only make an update to the path at the next network event or re-optimization. The default value is no pce-control.
Default	no pce-control

pce-report

Syntax	pce-report {enable disable inherit}
Context	config>router>mpls>lsp
Description	This command configures the reporting mode to a PCE of a LSP in a PCC. 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/disable PCE reporting for all SR-TE LSPs for the purpose of LSP database synchronization. This configuration will be inherited by all LSPs of a given type. The PCC will report both CSPF and non-CSPF LSP. The default value is disabled (no pce-report). This default value is meant to control the introduction of PCE into an existing network and let the operator decide if all LSP types need to be reported. The LSP level pce-report command overrides the global configuration for the reporting of LSP to PCE. The default value is to inherit the global MPLS level value. The inherit value returns the LSP to inherit the global configuration for that LSP type. Note that 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.
Default	pce-report inherit
Parameters	enable — Enables PCE reporting disable — Disables PCE reporting inherit — Inherits the global configuration for PCE reporting

priority

Syntax	priority setup-priority hold-priority no priority
Context	config>router>mpls>lsp>primary config>router>mpls>lsp>secondary
Description	This command enables the soft pre-emption procedures for this LSP path. The operator enables the soft pre-emption mechanism on a specific LSP name by explicitly configuring the setup and holding priorities for the primary path at the head-end node. The operator can similarly configure priority values for a secondary path for this LSP name. Different values could be used for the primary and for any of the secondary paths. In the absence of explicit user configuration, the setup priority is internally set to the default value of 7 and the holding priority is set to the default value of 0.



Note: Valid user-entered values for these two parameters require that the holding priority be numerically lower than or equal to the setup priority, otherwise pre-emption loops can occur.

Pre-emption is effected when a router pre-empting node processes a new RSVP session reservation and there is not enough available bandwidth on the RSVP interface, or the Class Type (CT) when Diff-Serv is enabled, to satisfy the bandwidth in the Flowspec object while there exist other session reservations for LSP paths with a strictly lower holding priority (numerically higher holding priority value) than the setup priority of the new LSP reservation. If enough available bandwidth is freed on the link or CT to accommodate the new reservation by pre-empting one or more lower priority LSP paths, the pre-empting node allows temporary overbooking of the RSVP interface and honors the new reservation.

The pre-empting node will immediately set the 'Preemption pending' flag (0x10) in the IPv4 Sub-Object in the RRO object in the Resv refresh for each of the pre-empted LSP paths. The IPv4 Sub-Object corresponds to the outgoing interface being used by the pre-empting and pre-empted LSP paths; however, the bandwidth value in the Flowspec object is not changed. The Resv flag must also be set if the pre-empting node is a merge point for the primary LSP path and the backup bypass LSP or detour LSP and the backup LSP is activated.

When evaluating if enough available bandwidth will be freed, the pre-empting node considers the reservations in order from the lowest holding priority (numerically higher holding priority value) to the holding priority just below the setup priority of the new reservation. A new reservation cannot pre-empt a reservation which has a value of the holding priority equal to the new reservation setup priority.

When Diff-Serv is enabled on the pre-empting node and the MAM bandwidth allocation model is used, a new reservation can only pre-empt a reservation in the same Class Type (CT).

LSP paths which were not flagged at the head-end for soft pre-emption will be hard pre-empted. LSP paths with the default holding priority of 0 cannot be pre-empted. LSP paths with zero bandwidth do not pre-empt other LSP paths regardless of the values of the path setup priority and the path holding priority. They can also not be pre-empted.

When evaluating if enough available bandwidth will be freed, the pre-empting node considers the reservations in order from the lowest holding priority (numerically higher holding priority) to the holding priority just below the setup priority of the new reservation. There is no specific order in which the reservations in the same holding priority are considered.

The pre-empting node starts a preemption timer for each of the pre-empted LSP paths. While this timer is on, the node should continue to refresh the Path and Resv for the pre-empted LSP paths. When the preemption timer expires, the node tears down the reservation if the head-end node has not already done so.

A head-end node upon receipt of the Resv refresh message with the 'Preemption pending' flag must immediately perform a make-before-break on the affected adaptive CSPF LSP. Both IGP metric and TE metric based CSPF LSPs are included. If an alternative path that excludes the flagged interface is not found, then the LSP is put on a retry in a similar way to the Global Revertive procedure at a head-end node. However, the number of retries and the retry timer are governed by the values of the retry-limit and retry-timer parameters:
config>router>mpls>lsp>retry-limit; config>router>mpls>lsp>retry-timer.

MPLS will keep the address list of flagged interfaces for a maximum of 60 s (not user-configurable) from the time the first Resv message with the 'Preemption pending' flag is received. This actually means that MPLS will request CSPF to find a path that excludes the flagged interfaces in the first few retries until success or until 60 s have elapsed. Subsequent retries after the 60 s will not exclude the flagged interfaces as it is assumed IGP has converged by then and the Unreserved Bandwidth sub-TLV for that priority, or TE Class, in the TE database will show the updated value taking into account the pre-empting LSP path reservation or a value of zero if overbooked.

If the LSP has a configured secondary standby which is operationally UP, the router will switch the path of the LSP to it and then start the MBB. If no standby path is available and a secondary non-standby is configured, the router will start the MBB and signal the path of the secondary. The LSP path will be switched to either the secondary or the new primary, whichever comes up first.

The no form of the command reverts the LSP path priority to the default values and results in setting the setup priority to 7, in setting the holding priority to 0, and in clearing the 'soft preemption desired' flag in the RRO in the Resv refresh message.

Default	no priority
Parameters	<i>setup-priority</i> — 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 pre-empt any other session. <i>holding-priority</i> — The priority of the reservation for this session at pre-emption action. Values 0 to 7 (0 is the highest priority and 7 is the lowest priority.) Default 0 — This session does not get pre-empted by any other session.

main-ct-retry-limit

Syntax	main-ct-retry-limit number no main-ct-retry-limit
Context	config>router>mpls>lsp
Description	This command configures the maximum number of retries the LSP primary path should be retried with the LSP Diff-Serv main Class Type (CT).

When an unmapped LSP primary path goes into retry, it uses the main CT until the number of retries reaches the value of the new main-ct-retry-limit parameter. If the path did not come up, it must start using the backup CT at that point in time. By default, this parameter is set to infinite value. The new main-ct-retry-limit parameter has no effect on an LSP primary path which retries due to a failure event.

An unmapped LSP primary path is a path which has never received a Resv in response to the first Path message sent. This can occur when performing a “shut/no-shut” on the LSP or LSP primary path or when the node reboots. An unmapped LSP primary path goes into retry if the retry timer expired or the head-end node received a PathErr message before the retry timer expired.

If the user entered a value of the main-ct-retry-limit parameter that is greater than the value of the LSP retry-limit, the number of retries will still stop when the LSP primary path reaches the value of the LSP retry-limit. In other words, the meaning of the LSP retry-limit parameter is not changed and always represents the upper bound on the number of retries. The unmapped LSP primary path behavior applies to both CSPF and non-CSPF LSPs.

The **no** form of this command sets the parameter to the default value of zero (0) which means the LSP primary path will retry forever.

Default	no main-ct-retry-limit
Parameters	<i>number</i> — The number of times MPLS will attempt to re-establish the LSP primary path using the Diff-Serv main CT. Allowed values are integers in the range of zero (0) to 10,000, where zero indicates to retry infinitely.
Values	0 to 1000, integer

max-sr-labels

Syntax	max-sr-labels <i>label-stack-size</i> [additional-frr-labels <i>labels</i>] no max-sr-labels
Context	config>router>mpls>lsp
Description	This command configures the maximum number of labels which the ingress LER can push for a given SR-TE LSP. This command is used to allow room to insert additional transport, service, and other labels when packets are forwarded in a given context. The max-sr-labels <i>label-stack-size</i> value should reflect the desired maximum label stack of the primary path of the SR-TE LSP. The value in additional-frr-labels <i>labels</i> 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 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.

The maximum label stack supported by the router is always signaled by PCC in the PCEP Open object as part of the as 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 PCC to PCE in the Segment-ID (SID) Depth value in a METRIC object for both a PCE computed LSP and a PCE controlled LSP. 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 Segment-ID (SID) Depth value if signaled, a reply with no path will be returned to PCC.

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

Parameters	<i>label-stack-size</i> — Specifies the label stack size.
Values	1 to 11
Default	6
	additional-frr-labels <i>labels</i> — Configures
Values	0 to 3
Default	1

metric

Syntax	[no] metric <i>metric</i>
Context	config>router>mpls>lsp config>router>mpls>lsp-template config>router>mpls>static-lsp
Description	This command allows the user to override the LSP operational metric with a constant administrative value that will not change regardless of the actual path the LSP is using over its lifetime. The LSP operational metric will match the metric the active path of this LSP is using at any given time. For a CSPF LSP, this metric represents the cumulative IGP metric of all the links the active path is using. If CSPF for this LSP is configured to use the TE metric, the LSP operational metric is set to the maximum value. For a non-CSPF LSP, the operational metric is the shortest IGP cost to the destination of the LSP.

The LSP operational metric is used by some applications to select an LSP among a set of LSPs that are destined to the same egress router. The LSP with the lowest operational metric will be selected. If more than one LSP with the same lowest LSP metric exists, the LSP with the lowest tunnel index will be selected. The configuration of a constant metric by the user will make sure the LSP always maintains its preference in this selection regardless of the path it is using at any given time. Applications that use the LSP operational metric include LDP-over-RSVP, VPRN auto-bind, and IGP, BGP and static route shortcuts.

The **no** form of this command disables the administrative LSP metric and reverts to the default setting in which the metric value will represent the LSP metric returned by MPLS. The same behavior is obtained if the user entered a metric of value zero (0).

Default no metric. The LSP operational metric defaults to the metric returned by MPLS.

Parameters *metric* — Specifies the integer value which specifies the value of the LSP administrative metric. A value of zero command reverts to the default setting and disables the administrative LSP metric.

Values 0 to 16777215

to

Syntax **to** [*ip-address* | *node-id* [*a.b.c.d* | 1...4294967295]]

Context config>router>mpls>lsp

Description This command specifies the system IP address or MPLS-TP node-id 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.

For a non MPLS-TP LSP, the **to** *ip-address* *must* be the system IP address of the egress router. If the **to** address does not match the SDP address, the LSP is not included in the SDP definition.

For an MPLS-TP LSP, the **to** *node-id* may be either in 4-octet IPv4 address format, or a 32-bit unsigned integer. This command is mandatory to create an MPLS-TP LSP. A value of zero is invalid. This **to** address is used in the MPLS-TP LSP ID, and the MPLS-TP MEP ID for the LSP.

Default no default

Parameters *ip-address* — The system IP address of the egress router.

node-id *a.b.c.d* | 1...4294967295 — 4-octet IPv4 formatted or unsigned 32-bit integer
MPLS-TP node-id of the egress router.

propagate-admin-group

Syntax	[no] propagate-admin-group
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command enables propagation of session attribute object with resource affinity (C-type 1) in PATH message. If a session attribute with resource affinity is received at an LSR, then it will check the compatibility of admin-groups received in PATH message against configured admin-groups on the egress interface of LSP. To support admin-group for inter-area LSP, the ingress node must configure propagating admin-groups within the session attribute object. If a PATH message is received by an LSR node that has the cspf-on-loose 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 ingress node. If this option is disabled, then the session attribute object without resource affinity (C-Type 7) is propagated in PATH message and CSPF at the LSR node will not include admin-group constraints. This admin group propagation is supported with a P2P LSP, a P2MP LSP instance, and an LSP template. The user can change the value of the propagate-admin-group option on the fly. A RSVP P2P LSP will perform a Make-Before-Break (MBB) on changing the configuration. A S2L path of an RSVP P2MP LSP will perform a Break-Before-Make on changing the configuration.
Default	no propagate-admin-group

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 names LSP can be used or no as part of the auto-bind feature for VPRN services. By default a names LSP is available for inclusion to used for the auto-bind feature. By configuring the command vprn-auto-bind exclude, the associated LSP will not be used by the auto-bind feature within VPRN services. The no form of the command resets the flag backto the default value.
Default	include
Parameters	include — Allows an associated LSPto be used by auto-bin for vprn services

exclude — Disables the use of the associated LSP to be used with the auto-bind feature for VPRN services.

retry-limit

Syntax	retry-limit <i>number</i> 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 LSP. 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 <i>lsp-name</i> no shutdown command to bring up the path after the retry-limit is exceeded. For P2MP LSP that are created based on the LSP template, all S2Ls must attempt to retry-limit before the client application is informed of failure. The no form of this command revert the parameter to the default value.
Default	0 (no limit, retries forever)
Parameters	<i>number</i> — The number of times 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 <i>seconds</i> no retry-timer
Context	config>router>mpls>lsp config>router>mpls>lsp-template
Description	This command configures the time (in s), for LSP re-establishment attempts after it has failed. The retry time is jittered to +/- 25% of its nominal value. For P2MP LSP created based on LSP template, all S2Ls must attempt to retry-limit before client application is informed of failure. The no form of this command reverts to the default value.

Default	30
Parameters	<i>seconds</i> — The amount of time (in s), between attempts to re-establish the LSP after it has failed. Allowed values are integers in the range of 1 to 600.
Values	1 to 600

revert-timer

Syntax	revert-timer <i>timer-value</i> no revert-timer
Context	config>router>mpls>lsp
Description	This command configures a revert timer on an LSP. The timer starts when the LSP primary path recovers from a failure. The LSP reverts from a secondary path to the primary path when the timer expires, or when the secondary path fails. The no form of this command cancels any currently outstanding revert timer. If the LSP is up when a no revert-timer is issued, the LSP will revert to the primary path. Otherwise the LSP reverts when the primary path is restored.
Default	no revert-timer
Parameters	<i>timer-value</i> — The amount of time, in one minute increments, between attempts to re-establish the LSP after it has failed. Allowed values are integers in the range of 0 to 4320. Values 0 to 4320

rsvp-resv-style

Syntax	rsvp-resv-style [<i>se</i> <i>ff</i>]
Context	config>router>mpls>lsp
Description	This command specifies the RSVP reservation style, shared explicit (se) or fixed filter (ff). A reservation style is a set of control options that specify a number of supported parameters. The style information is part of the LSP configuration.
Default	se
Parameters	<i>ff</i> — 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-template
Description	This command disables the existing LSP including the primary and any standby secondary paths. To shutdown only the primary enter the config router mpls lsp lsp-name primary path-name shutdown command. To shutdown a specific standby secondary enter the config router mpls lsp lsp-name secondary path-name shutdown command. The existing configuration of the LSP is preserved. Use the no form of this command to restart the LSP. LSPs are created in a shutdown state. Use this command to administratively bring up the LSP.
Default	shutdown

lsp-template

Syntax	[no] lsp-template <i>lsp-template-name</i> p2mp-lsp
Context	config>router>mpls
Description	This command creates a template construct that can be referenced by client application where dynamic LSP creation is required. 'p2mp-lsp' keyword is mandatory. The no form of command deletes LSP template. LSP template cannot be deleted if a client application is using it.
Default	none
Parameters	<i>lsp-template-name</i> — Name to identify LSP template. Any LSP template name and LSP name must not be same.

default-path

Syntax	[no] default-path <i>path-name</i>
Context	config>router>mpls>lsp-template
Description	A default path binding must be provided before the LSP template can be used for signaling LSP. The LSP template must be shutdown to modify default-path binding. The no form of command will delete path binding.
Default	none
Parameters	<i>path-name</i> — Configures the default path binding

lsp-bfd

Syntax	lsp-bfd no lsp-bfd
Context	config>router
Description	This command creates a context for the configuration of LSP BFD parameters.
Default	no lsp-bfd

bfd-sessions

Syntax	bfd-sessions <i>max-limit</i> no bfd-sessions
Context	config>router>lsp-bfd
Description	This command enables or disables LSP BFD at the tail end of LSPs on the system. It is also used to limit the maximum number of LSP BFD sessions that may be established at the tail-end of LSPs on a node to <i>max-limit</i> . It has no impact on the number of LSP BFD sessions that may be configured at the head end.
Default	no bfd-sessions : The establishment of LSP BFD sessions by the node at the tail end of LSPs is disabled.
Parameters	<i>max-limit</i> — The maximum number of LSP BFD sessions at the tail end of LSPs that can be established on a system. The maximum value that can be entered is constrained by the system wide limit for centralized BDF sessions. Values 1- max, where max is the platform specific limit on centralized BFD sessions.

Primary and Secondary Path Commands

primary

Syntax	primary <i>path-name</i> no primary
Context	config>router>mpls>lsp
Description	This command specifies a preferred path for the LSP. This command is optional only if the secondary <i>path-name</i> 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 <i>path-name</i> command, override the LSP attributes. The no form of this command deletes the association of this <i>path-name</i> from the LSP <i>lsp-name</i>. All configurations specific to this primary path, such as record, bandwidth, and hop limit, are deleted. The primary path must be shutdown 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	none
Parameters	<i>path-name</i> — The case-sensitive alphanumeric name label for the LSP path up to 32 characters in length.

secondary

Syntax	[no] secondary <i>path-name</i>
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 <i>lsp-name</i> primary <i>path-name</i> command is specified. After the switch over from the primary to the secondary, the software continuously tries to revert to the primary path. The switch back to the primary path is based on the retry-timer interval. Up to eight secondary paths can be specified. All the secondary paths are considered equal and the first available path is used. The software will not switch back among secondary paths. Software starts the signaling of all non-standby secondary paths at the same time. Retry counters are maintained for each unsuccessful attempt. Once the retry limit is reached on a path, software will not attempt to signal the path and administratively shuts down the path. The first successfully established path is made the active path for the LSP.

The **no** form of this command removes the association between this *path-name* and *lsp-name*. All specific configurations for this association are deleted. The secondary path must be shutdown 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	none
Parameters	<i>path-name</i> — 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 functionality for an LSP or a primary or secondary LSP path. When enabled for the LSP, make-before-break will be performed for primary path and all the secondary paths of the LSP.
Default	adaptive

backup-class-type

Syntax	backup-class-type ct-number no backup-class-type
Context	config>router>mpls>lsp>primary
Description	This command enables the use of the Diff-Serv backup Class-Type (CT), instead of the Diff-Serv main CT, to signal the LSP primary path when it fails and goes into retry. The Diff-Serv main CT is configured at the LSP level or at the primary path level using the following commands:

```
config>router>mpls>lsp>class-type ct-number
```

```
config>router>mpls>lsp>primary>class-type ct-number
```

When a LSP primary path retries due a failure, for example, it fails after being in the UP state, or undergoes any type of Make-Before-Break (MBB), MPLS will retry a new path for the LSP using the main CT. If the first attempt failed, the head-end node performs subsequent retries using the backup CT. This procedure must be followed regardless if the currently used CT by this path is the main or backup CT. This applies to both CSPF and non-CSPF LSPs.

The triggers for using the backup CT after the first retry attempt are:

1. A local interface failure or a control plane failure (hello timeout etc.).
2. Receipt of a PathErr message with a notification of a FRR protection becoming active downstream and/or Receipt of a Resv message with a 'Local-Protection-In-Use' flag set. This invokes the FRR Global Revertive MBB.
3. Receipt of a PathErr message with error code=25 ("Notify") and sub-code=7 ("Local link maintenance required") or a sub-code=8 ("Local node maintenance required"). This invokes the TE Graceful Shutdown MBB.
4. Receipt of a Resv refresh message with the 'Preemption pending' flag set or a PathErr message with error code=34 ("Reroute") and a value=1 ("Reroute request soft preemption"). This invokes the soft preemption MBB.
5. Receipt of a ResvTear message.
6. A configuration change MBB.
7. The user executing the clear>router>mpls>lsp command.

When an unmapped LSP primary path goes into retry, it uses the main CT until the number of retries reaches the value of the new **main-ct-retry-limit** parameter. If the path did not come up, it must start using the backup CT at that point in time. By default, this parameter is set to infinite value. The new main-ct-retry-limit parameter has no effect on an LSP primary path which retries due to a failure event.

An unmapped LSP primary path is a path which has never received a Resv in response to the first Path message sent. This can occur when performing a 'shut/no-shut' on the LSP or LSP primary path or when the node reboots. An unmapped LSP primary path goes into retry if the retry timer expired or the head-end node received a PathErr message before the retry timer expired.

When the re-signal timer expires, CSPF will try to find a path with the main CT. The head-end node must re-signal the LSP even if the new path found by CSPF is identical to the existing one since the idea is to restore the main CT for the primary path. A path with main CT is not found, the LSP remains on its current primary path using the backup CT.

When the user performs a manual re-signal of the primary path, CSPF will try to find a path with the main CT. The head-end node must re-signal the LSP as in current implementation.

The **no** form of this command disables the use of the Diff-Serv backup CT.

Default	no backup-class-type
Parameters	<i>ct-number</i> — The Diff-Serv Class Type number. One or more system forwarding classes can be mapped to a CT.
Values	0-7, integer

bandwidth

Syntax	bandwidth <i>rate-in-mbps</i> no bandwidth
Context	config>router>mpls>lsp>primary config>router>mpls>lsp>secondary config>router>mpls>lsp>template>fast-reroute
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	<i>rate-in-mbps</i> — The amount of bandwidth reserved for the LSP path in Mb/s. Allowed values are integers in the range of 1 to 100000. Values 0 to 100000

exclude

Syntax	[no] exclude <i>group-name</i> [<i>group-name...</i> (up to 5 max)]
Context	config>router>mpls>lsp config>router>mpls>lsp>primary config>router>mpls>lsp>primary-p2mp-instance config>router>mpls>lsp>secondary
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>admin-group context. The config>router>mpls>lsp>primary-p2mp-instance>exclude command is not supported on the 7450 ESS. Use the no form of the command to remove the exclude command.
Default	no exclude
Parameters	<i>group-name</i> — Specifies the existing group-name to be excluded when an LSP is set up.

hop-limit

Syntax	hop-limit <i>number</i> 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 with the following implications: If the new value is less than the current hops of the established LSP, the LSP is brought down. MPLS then tries to re-establish the LSP within the new hop-limit number. If the new value is equal or more than the current hops of the established LSP then the LSP will be unaffected. The no form of this command reverts the values defined under the LSP definition using the config router mpls lsp lsp-name hop-limit command.
Default	no hop-limit
Parameters	<i>number</i> — The number of hops the LSP can traverse, expressed as an integer. Values 2 to 255

record

Syntax	[no] record
Context	config>router>mpls>lsp>primary config>router>mpls>lsp>primary-p2mp-instance config>router>mpls>lsp>secondary config>router>mpls>lsp>template
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 config>router>mpls>lsp>primary-p2mp-instance>record command is not supported on the 7450 ESS. 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
config>router>mpls>lsp>template
```

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 computation of a secondary path for an LSP at the head-end LER.

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

If CSPF is not enabled on the LSP name, then a secondary path of that LSP which has the SRLG constraint included will be shut down and a specific failure code will indicate the exact reason for the failure in **show>router>mpls>lsp>path>detail** output.

At initial primary LSP path establishment, if primary does not come up or primary is not configured, SRLG secondary will not be signaled and will put to down state. A specific failure code will indicate the exact reason for the failure in **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 task will signal it and the LSP use it.

As soon as the primary path is configured and successfully established, MPLS/RSVP moves the LSP to the primary and signals all SRLG secondary paths.

Any time the primary path is re-optimized, has undergone MBB, or has come back up after being down, MPLS/RSVP task checks with CSPF if the SRLG secondary should be re-signalized. If MPLS/RSVP finds that current secondary path is no longer SRLG disjoint, for example, it became ineligible, it puts it on a delayed MBB immediately after the expiry of the retry timer. If MBB fails at the first try, the secondary path is torn down and the path is put on retry.

At the next opportunity the primary goes down, the LSP will use the path of an eligible SRLG secondary if it is UP. If all secondary eligible SLRG paths are Down, MPLS/RSVP will use a non SRLG secondary if configured and UP. If while the LSP is using a non SRLG secondary, an eligible SRLG secondary came back up, MPLS/RSVP will not switch the path of the LSP to it. As soon as primary is re-signaled and comes up with a new SLRG list, MPLS/RSVP will re-signal the secondary using the new SRLG list.

A secondary path which becomes ineligible as a result of an update to the SRLG membership list of the primary path will have the ineligibility status removed on any of the following events:

1. A successful MBB of the standby SRLG path which makes it eligible again.
2. The standby path goes down. MPLS/RSVP puts the standby on retry at the expiry of the retry timer. If successful, it becomes eligible. If not successful after the retry-timer expired or the number of retries reached the number configured under the retry-limit parameter, it is left down.
3. The primary path goes down. In this case, the ineligible secondary path is immediately torn down and will only be re-signaled when the primary comes back up with a new SRLG list.

Once primary path of the LSP is setup and is operationally up, any subsequent changes to the SRLG group membership of an interface the primary path is using would not be considered until the next opportunity the primary path is re-signaled. The primary path may be re-signaled due to a failure or to a make-before-break operation. Make-before-break occurs as a result of a global revertive operation, a timer based or manual re-optimization of the LSP path, or an operator change to any of the path constraints.

One an SRLG secondary path is setup and is operationally UP, any subsequent changes to the SRLG group membership of an interface the secondary path is using would not be considered until the next opportunity secondary path is re-signaled. The secondary path is re-signaled due to a failure, to a re-signaling of the primary path, or to a make before break operation. Make-before break occurs as a result of a timer based or manual re-optimization of the secondary path, or an operator change to any of the path constraints of the secondary path, including enabling or disabling the SRLG constraint itself.

Also, the user-configured include/exclude admin group statements for this secondary path are also checked together with the SRLG constraints by CSPF. Finally, enabling SRPG on a secondary standby path that is in the up state will cause the path to be torn down and re-signaled using the SRLG constraint.

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 once the primary path LSP fails. The standby keyword ensures that the secondary path LSP is signaled and maintained indefinitely in a hot-standby state. When the primary path is re-established then the traffic is switched back to the primary path LSP. The no form of this command specifies that the secondary LSP is signaled when the primary path LSP fails.
Default	none

path-preference

Syntax	[no] path-preference <i>value</i>
Context	config>router>mpls>lsp>secondary
Description	This command enables use of path preference among configured standby secondary paths per LSP. If all standby secondary paths have a default path-preference value then a non-standby secondary path will remain the active path while a standby secondary is available. A standby secondary path configured with highest priority (lowest path-preference value) must be made the active path when the primary is not in use. Path preference can be configured on standby secondary path. The no form of this command resets the path-preference to the default value.
Default	255
Parameters	<i>value</i> — Specifies an alternate path for the LSP if the primary path is not available, 1–255

LSP Path Commands

hop

Syntax	hop <i>hop-index ip-address</i> { strict loose } no hop <i>hop-index</i>
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 or the system IP address. If the system IP address is specified then the LSP can choose the best available interface.

Optionally, the LSP ingress and egress IP address can be included as the first and the 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.

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 shutdown first in order to delete the hop from the hop list. The **no hop hop-index** command will not result in any action except a warning message on the console indicating that the path is administratively up.

Default	none
Parameters	<i>hop-index</i> — The hop index is used to order the hops specified. The LSP always traverses from the lowest hop index to the highest. The hop index does not need to be sequential. Values 1 to 1024 <i>ip-address</i> — The system or network interface IP address of the transit router. The IP address can be the interface IP address or the system IP address. If the system IP address is specified then the LSP can choose the best available interface. A hop list can also include the ingress interface IP address, the system IP address, and the egress IP address of any of the specified hops. loose — This keyword specifies that the route taken by the LSP from the previous hop to this hop can traverse through other routers. Multiple hop entries with the same IP address are flagged as errors. Either the loose or strict keyword must be specified. strict — This keyword 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, then that is the interface the LSP must use. If there are direct parallel links between the previous router and this router and if 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. Multiple hop entries with the same IP address are flagged as errors. Either the loose or strict keyword must be specified.

path

Syntax	[no] path <i>path-name</i>
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 <i>path-name</i> specified) in which case the LSP is set up based on IGP (best effort) calculated shortest path to the egress router. Paths are created in a shutdown state. A path must be shutdown before making any changes (adding or deleting hops) to the path. When a path is shutdown, 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 **shutdown** and unbound from all LSPs using the path before it can be deleted. The **no path path-name** command will not result in any action except a warning message on the console indicating that the path may be in use.

Parameters *path-name* — Specify a unique case-sensitive alphanumeric name label for the LSP path up to 32 characters in length.

shutdown

Syntax **[no] shutdown**

Context config>router>mpls>path

Description This command disables the existing LSPs using this path. All services using these LSPs are affected. Binding information, however, is retained in those LSPs. Paths are created in the **shutdown** state.

The **no** form of this command administratively enables the path. All LSPs, where this path is defined as primary or defined as standby secondary, are (re)established.

Default **shutdown**

Static LSP Commands

static-lsp

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

Context config>router>mpls

Description This command is used to configure a static LSP on the ingress router. The static LSP is a manually set up LSP where the nexthop 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 shutdown first in order to delete it. If the LSP is not shut down, the **no static-lsp lsp-name** command does nothing except generate a warning message on the console indicating that the LSP is administratively up.

Parameters *lsp-name* — Name that identifies the LSP.
 Values Up to 32 alphanumeric characters.

static-lsp-fast-retry

Syntax **static-lsp-fast-retry** *seconds*
 no static-lsp-fast-retry

Context config>router>mpls

Description This command specifies the value used as the fast retry timer for a static LSP.

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

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

Default no static-fast-retry-timer

Parameters *seconds* — specifies the value (in s), used as the fast retry timer for a static LSP.
 Values 1-30

metric

Syntax **metric** *metric*
 no metric

Context config>router>mpls>static-lsp

Description This command configures the MPLS static LSP metric.

Parameters *metric* — Specifies the static LSP metric.
 Values 0 to 16777215

push

Syntax **push** {*label* | **implicit-null-label**} **nexthop** *ip-address*
 no push {*out-label* | **implicit-null-label**}

Context config>router>mpls>static-lsp

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

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

Parameters *implicit-null-label* — Specifies the use of the implicit label value for the push operation.

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

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

Values 16 to 1048575

nexthop *ip-address* — This command specifies the IP address of the next hop towards the LSP egress router. If an ARP entry for the next hop exists, then the static LSP is marked operational. If ARP entry does not exist, software sets the operational status of the static LSP to down and continues to ARP for the configured nexthop. Software continuously tries to ARP for the configured nexthop at a fixed interval.

shutdown

Syntax [no] shutdown

Context config>router>mpls>static-lsp

Description This command is used to administratively disable the static LSP.

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

Default shutdown

to

Syntax to *ip-address*

Context config>router>mpls>static-lsp

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

Default none

Parameters *ip-address* — The system IP address of the egress router.

2.20.2.1.3 Point-to-Multipoint MPLS (P2MP) Commands

p2mp-id

Syntax	p2mp-id <i>id</i>
Context	config>router>mpls>lsp
Description	This command configures the identifier of an RSVP P2MP LSP. An RSVP P2MP LSP is fully identified by the combination of: <P2MP ID, tunnel ID, extended tunnel ID> part of the P2MP session object, and <tunnel sender address, LSP ID> fields in the p2mp sender_template object. The p2mp-id is a 32-bit identifier used in the session object that remains constant over the life of the P2MP tunnel. It is unique within the scope of the ingress LER. The no form restores the default value of this parameter. This command is not supported on the 7450 ESS.
Default	0
Parameters	<i>id</i> — Specifies a P2MP identifier.
Values	0 to 65535

primary-p2mp-instance

Syntax	[no] primary-p2mp-instance <i>instance-name</i>
Context	config>router>mpls>lsp
Description	This command creates the primary instance of a P2MP LSP. The primary instance of a P2MP LSP is modeled as a set of root-to-leaf (S2L) sub-LSP's. The root, for example a head-end node triggers signaling using one path message per S2L path. The leaf sub-LSP paths are merged at branching points. This command is not supported on the 7450 ESS.
Default	none
Parameters	<i>instance-name</i> — Specifies a name that identifies the P2MP LSP instance. The instance name can be up to 32 characters long and must be unique.

s2l-path

Syntax	[no] s2l-path <i>path-name to ip-address</i>

Context	config>router>mpls>lsp>primary-p2mp-instance
Description	This command creates a root-to-leaf (S2L) sub-LSP path for the primary instance of a P2MP LSP. The primary instance of a P2MP LSP is modeled as a set of root-to-leaf (S2L) sub-LSPs. The root, for example, head-end node, triggers signaling using one path message per S2L path. The leaf sub-LSP paths are merged at branching points. Each S2L sub-LSP is signaled in a separate path message. Each leaf node will respond with its own RESV message. A branch LSR node will forward the path message of each S2L sub-LSP to the downstream LSR without replicating it. It will also forward the RESV message of each S2L sub-LSP to the upstream LSR without merging it with the RESV messages of other S2L sub-LSPs of the same P2MP LSP. The same is done for subsequent refreshes of the path and RESV states. The S2L paths can be empty paths or can specify a list of explicit hops. The path name must exist and must have been defined using the config>router>mpls>path command. The same path name can be re-used by more than one S2L of the primary P2MP instance. However, the to keyword must have a unique argument per S2L as it corresponds to the address of the egress LER node. This command is not supported on the 7450 ESS.
Default	none
Parameters	<i>path-name</i> — Specifies the name of the path which consists of up to 32 alphanumeric characters. to <i>ip-address</i> — Specifies the system IP address of the egress router.

p2mp-resignal-timer

Syntax	p2mp-resignal-timer <i>minutes</i> no p2mp-resignal-timer
Context	config>router>mpls
Description	This command configures the re-signal timer for a P2MP LSP instance. MPLS will request CSPF to re-compute the whole set of S2L paths of a given active P2MP instance each time the P2MP re-signal timer expires. The P2MP re-signal timer is configured separately from the P2P LSP parameter. MPLS performs a global MBB and moves each S2L sub-LSP in the instance into its new path using a new P2MP LSP ID if the global MBB is successful, regardless of the cost of the new S2L path. This command is supported on the 7750 SR, 7950 XRS, and with VPLS only on the 7450 ESS. The no form of this command disables the timer-based re-signaling of P2MP LSPs on this system.

Parameters	<i>minutes</i> — Specifies the time MPLS waits before attempting to re-signal the P2MP LSP instance.
Values	60 to 10080

2.20.2.2 RSVP Commands

2.20.2.2.1 Generic Commands

shutdown

Syntax	[no] shutdown
Context	config>router>rsvp config>router>rsvp>interface
Description	This command disables the RSVP protocol instance or the RSVP-related functions for the interface. The RSVP configuration information associated with this interface is retained. When RSVP is administratively disabled, all the RSVP sessions are torn down. The existing configuration is retained. The no form of this command administratively enables RSVP on the interface.
Default	shutdown

2.20.2.2.2 RSVP Commands

rsvp

Syntax	[no] rsvp
Context	config>router
Description	This command enables the context to configure RSVP protocol parameters. RSVP is not enabled by default and must be explicitly enabled (no shutdown). RSVP is used to set up LSPs. RSVP should be enabled on all router interfaces that participate in signaled LSPs.

The **no** form of this command deletes this RSVP protocol instance and removes all configuration parameters for this RSVP instance. To suspend the execution and maintain the existing configuration, use the **shutdown** command. RSVP must be shutdown before the RSVP instance can be deleted. If RSVP is not shutdown, the **no rsvp** command does nothing except issue a warning message on the console indicating that RSVP is still administratively enabled.

Default no shutdown

entropy-label-capability

Syntax [no] **entropy-label-capability**

Context config>router>rsvp
config>router>ldp

Description This command enables or disables ELC for RSV.

If **entropy-label-capability** is configured, then the system will signal (using the procedures specified in RFC 6790) that it is capable of receiving and processing the entropy label and ELI on incoming packets of RSVP and LDP LSPs.

If **no entropy-label-capability** is configured, then the system will not signal ELC. If an ELI is exposed on a packet where the tunnel label is popped at the termination of that LSP, and an entropy label is not configured, then the packet will be dropped.

Default no entropy-label-capability

diffserv-te

Syntax **diffserv-te** [mam | rdm]
no diffserv-te

Context config>router>rsvp

Description This command enabled Diff-Serv Traffic Engineering on the node.

When this command is enabled, IS-IS and OSPF will start advertising available bandwidth for each TE class configured under the diffserv-te node. This command will only have effect if the operator has already enabled traffic engineering at the IS-IS and/or OSPF routing protocol levels:

config>router>isis>traffic-engineering

and/or:

config>router>ospf>traffic-engineering

IGP will advertize for each RSVP interface in the system the available bandwidth in each TE class in the unreserved bandwidth TE parameter for that class. In addition, IGP will continue to advertize the existing Maximum Reservable Link Bandwidth TE parameter to mean the maximum bandwidth that can be booked on a given interface by all classes. The value advertized is adjusted with the link **subscription percentage** factor configured in the **config>router>rsvp>interface** context.

The user configures the following parameters for the operation of Diff-Serv:

- Definition of TE classes, TE Class = {Class Type (CT), LSP priority}.
- Mapping of the system forwarding classes to the Diff-Serv Class Type (CT).
- Configuration of the percentage of RSVP interface bandwidth each CT shares, i.e., the Bandwidth Constraint (BC).

When Diff-Serv TE is enabled, the system will automatically enable the Max Allocation Model (MAM) Admission Control Policy. MAM represents the bandwidth constraint model for the admission control of an LSP reservation to a link. This is the only Admission Control Policy supported in this release.

Each CT shares a percentage of the Maximum Reservable Link Bandwidth via the user configured Bandwidth Constraint (BC) for this CT. The Maximum Reservable Link Bandwidth is the link bandwidth multiplied by the RSVP interface subscription factor.

The sum of all BC values across all CTs will not exceed the Maximum Reservable Link Bandwidth. In other words, the following rule is enforced:

$$\text{SUM}(BC_c) \leq \text{Max-Reservable-Bandwidth}, 0 \leq c \leq 7$$

An LSP of class-type CT_c , setup priority p , holding priority h ($h \leq p$), and bandwidth B is admitted into a link if the following condition is satisfied:

$$B \leq \text{Unreserved Bandwidth for TE-Class}[i]$$

where TE-Class $[i]$ maps to $\langle CT_c, p \rangle$ in the definition of the TE classes on the node. The bandwidth reservation is effected at the holding priority, i.e., in TE-class $[j] = \langle CT_c, h \rangle$. Thus, the reserved bandwidth for CT_c and the unreserved bandwidth for the TE classes using CT_c are updated as follows:

$$\text{Reserved}(CT_c) = \text{Reserved}(CT_c) + B$$

$$\text{Unreserved TE-Class } [j] = BC_c - \text{SUM}(\text{Reserved}(CT_c, q)) \text{ for } 0 \leq q \leq h$$

$$\text{Unreserved TE-Class } [i] = BC_c - \text{SUM}(\text{Reserved}(CT_c, q)) \text{ for } 0 \leq q \leq p$$

The same is done to update the unreserved bandwidth for any other TE class making use of the same CT_c . These new values are advertised to the rest of the network at the next IGP-TE flooding.

The Russian Doll Model (RDM) LSP admission control policy allows bandwidth sharing across Class Types. It provides a hierarchical model by which the reserved bandwidth of a CT is the sum of the reserved bandwidths of the numerically equal and higher CTs.

The RDM model is defined using the following equations:

$$\text{SUM (Reserved (CTc))} \leq \text{BCb},$$

where the SUM is across all values of c in the range $b \leq c \leq (\text{MaxCT} - 1)$, and BCb is the bandwidth constraint of CTb.

BC0= Max-Reservable-Bandwidth, so that

$$\text{SUM (Reserved(CTc))} \leq \text{Max-Reservable-Bandwidth},$$

where the SUM is across all values of c in the range $0 \leq c \leq (\text{MaxCT} - 1)$.

When Diff-Serv is disabled on the node, this model degenerates into a single default CT internally with eight pre-emption priorities and a non-configurable BC equal to the Maximum Reservable Link Bandwidth. This would behave exactly like CT0 with eight pre-emption priorities and BC= Maximum Reservable Link Bandwidth if Diff-Serv was enabled.

The enabling or disabling of Diff-Serv TE on the system requires the RSVP and MPLS protocol be shutdown.

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

Default	no diffserv-te
Parameters	mam — Defines the default admission control policy for Diff-Serv LSPs. rdm — Defines Russian doll model for the admission control policy of Diff-Serv LSPs.

class-type-bw

Syntax	class-type-bw ct0 %link-bandwidth ct1%link-bandwidth ct2%link-bandwidth ct3%link-bandwidth ct4%link-bandwidth ct5%link-bandwidth ct6%link-bandwidth ct7%link-bandwidth no class-type-bw
Context	config>router>rsvp>diffserv-te config>router>rsvp>interface
Description	This command configures the percentage of RSVP interface bandwidth each CT shares, for example, the Bandwidth Constraint (BC). The absolute value of the CT share of the interface bandwidth is derived as the percentage of the bandwidth advertised by IGP in the Maximum Reservable Link Bandwidth TE parameter, for example, the link bandwidth multiplied by the RSVP interface subscription percentage parameter.



Note: This configuration also exists at RSVP interface level and the interface specific configured value overrides the global configured value. The BC value can be changed at any time.

The RSVP interface **subscription percentage** parameter is configured in the **config>router>rsvp>interface** context.

The operator can specify the Bandwidth Constraint (BC) for a CT which is not used in any of the TE class definition but that does not get used by any LSP originating or transiting this node.

When Diff-Serv is disabled on the node, this model degenerates into a single default CT internally with eight pre-emption priorities and a non configurable BC equal to the Maximum Reservable Link Bandwidth. This would behave exactly like CT0 with eight pre-emption priorities and BC= Maximum Reservable Link Bandwidth if Diff-Serv was enabled.

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

Parameters **ct0 (ct1/ct2/ — ct7) % link-bandwidth** — The Diff-Serv Class Type number. One or more system forwarding classes can be mapped to a CT.

Values 0 to 100 %

Default 0

fc

Syntax **fc fc-name class-type ct-number**
no fc fc-name

Context config>router>rsvp>diffserv-te

Description This command maps one or more system forwarding classes to a Diff-Serv Class Type (CT). The default mapping is shown in [Table 13](#).

Table 13 Forwarding Classes Mapping

FC ID	FC Name	FC Designation	Class Type (CT)
7	Network Control	NC	7
6	High-1	H1	6
5	Expedited	EF	5
4	High-2	H2	4
3	Low-1	L1	3

Table 13 Forwarding Classes Mapping (Continued)

FC ID	FC Name	FC Designation	Class Type (CT)
2	Assured	AF	2
1	Low-2	L2	1
0	Best Effort	BE	0

The **no** form of this command reverts to the default mapping for the forwarding class name.

Parameters **class-type** *ct-number* — The Diff-Serv Class Type number. One or more system forwarding classes can be mapped to a CT.

Values 0 to 7

te-class

Syntax **te-class** *te-class-number* **class-type** *ct-number* **priority** *priority*
no te-class *te-class-number*

Context config>router>rsvp>diffserv-te

Description This command configures a traffic engineering class. A TE class is defined as:

TE Class = {Class Type (CT), LSP priority}

Eight TE classes are supported. There is no default TE class once Diff-Serv is enabled. The user has to explicitly define each TE class.

When Diff-Serv is disabled, there will be an internal use of the default CT (CT0) and eight pre-emption priorities as shown in [Table 14](#).

Table 14 Default Class Type

Class Type (CT internal)	LSP Priority
0	7
0	6
0	5
0	4
0	3
0	2

Table 14 Default Class Type (Continued)

Class Type (CT internal)	LSP Priority
0	1
0	0

The **no** form of this command deletes the TE class.

- Parameters**
- te-class** *te-class-number* — The traffic engineering class number.
Values 0 to 7
- class-type** *ct-number* — The Diff-Serv Class Type number. One or more system forwarding classes can be mapped to a CT.
Values 0 to 7
- priority** *priority* — The LSP priority.
Values 0 to 7

gr-helper

- Syntax** **gr-helper** [**enable** | **disable**]
- Context** config>router>rsvp>if
- Description** This command enables the RSVP Graceful Restart Helper feature.

The RSVP-TE Graceful Restart helper mode allows the SR OS based system (the helper node) to provide another router that has requested it (the restarting node) a grace period, during which the system will continue to use RSVP sessions to neighbors requesting the grace period. This is typically used when another router is rebooting its control plane but its forwarding plane is expected to continue to forward traffic based on the previously available Path and Resv states.

The user can enable Graceful Restart helper on each RSVP interface separately. When the GR helper feature is enabled on an RSVP interface, the node starts inserting a new Restart_Cap Object in the Hello packets to its neighbor. The restarting node does the same and indicates to the helper node the desired Restart Time and Recovery Time.

The GR Restart helper consists of a couple of phases. Once it loses Hello communication with its neighbor, the helper node enters the Restart phase. During this phase, it preserves the state of all RSVP sessions to its neighbor and waits for a new Hello message.

Once the Hello message is received indicating the restarting node preserved state, the helper node enters the recovery phase in which it starts refreshing all the sessions that were preserved. The restarting node will activate all the stale sessions that are refreshed by the helper node. Any Path state which did not get a Resv message from the restarting node once the Recovery Phase time is over is considered to have expired and is deleted by the helper node causing the proper Path Tear generation downstream.

The duration of the restart phase (recovery phase) is equal to the minimum of the neighbor's advertised Restart Time (Recovery Time) in its last Hello message and the locally configured value of the max-restart (max-recovery) parameter.

When GR helper is enabled on an RSVP interface, its procedures apply to the state of both P2P and P2MP RSVP LSP to a neighbor over this interface.

Default disable

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 or all RSVP interfaces on the node if applied at the RSVP level. These are referred to as maintenance interface and maintenance node, respectively. To initiate a graceful shutdown the maintenance node generates a PathErr message with a specific error sub-code of Local Maintenance on TE Link required for each LSP that is exiting the maintenance interface. The node performs a single make-before-break attempt for all adaptive CSPF LSPs it originates and LSP paths using the maintenance interfaces. If an alternative path for an affected LSP is not found, then the LSP is maintained on its current path. The maintenance node also tears down and re-signals any detour LSP path using listed maintenance interfaces as soon as they are not active. The maintenance node floods an IGP TE LSA/LSP containing Link TLV for the links under graceful shutdown with Traffic Engineering metric set to 0xffffffff and Unreserved Bandwidth parameter set to zero (0). A head-end LER node, upon receipt of the PathErr message performs a single make-before-break attempt on the affected adaptive CSPF LSP. If an alternative path is not found, then 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:

- a. 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 which can be found.
- b. An adaptive CSPF LSP whose path has explicit hops defined using the listed maintenance interface(s)/node(s).
- c. A CSPF LSP with the adaptive option disabled and which current path is over the listed maintenance interfaces in the PathErr message. These are not subject to make-before-break.
- d. A non CSPF LSP which current path is over the listed maintenance interfaces in the PathErr message.

The head-end LER node upon receipt of the updates IPG TE LSA/LSP for the maintenance interfaces updates the TE database. This information will be used at the next scheduled CSPF computation for any LSP which path may 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 none

gr-helper-time

Syntax **gr-helper-time max-recovery** *recovery-interval* [1..1800] seconds **max-restart** *restart-interval*
no gr-helper-time

Context config>router>rsvp

Description This command configures the local values for the max-recovery and the max-restart intervals used in the RSVP Graceful Restart Helper feature.

The values are configured globally in RSVP but separate instances of the timers are applied to each RSVP interface that has the RSVP Graceful Restart Helper enabled.

The **no** version of this command re-instates the default value for the delay timer.

Parameters *recovery-interval* — Specifies the max recovery interval value in seconds.

Values 1 to 1800

Default 300

restart-interval — Specifies the max restart interval value in seconds.

Values 1 to 300

Default 120

implicit-null-label

Syntax	[no] implicit-null-label implicit-null-label
Context	config>router>rsvp
Description	This command enables the use of the implicit null label. Signalling the IMPLICIT NULL label value for all RSVP LSPs can be enabled for which this node is the egress LER. RSVP must be shutdown before being able to change this configuration option. The egress LER does not signal the implicit null label value on P2MP RSVP LSPs. However, the Penultimate Hop Popping (PHP) node can honor a resv message with the label value set to the implicit null. 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 <i>number</i> no keep-multiplier
Context	config>router>rsvp
Description	The keep-multiplier <i>number</i> is an integer used by RSVP to declare that a reservation is down or the neighbor is down. The no form of this command reverts to the default value.
Default	3
Parameters	<i>number</i> — The keep-multiplier value.
Values	1 to 255

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 PLR node or terminating on this 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 the latter will not accept summary refresh messages received over the bypass.

When disabled, the node as a PLR or MP will not set the “Refresh-Reduction-Capable” bit on RSVP messages pertaining to LSP paths tunneled over the bypass. It will also not send Message-ID in RSVP messages. This effectively disables summary refresh.

Default disable

rapid-retransmit-time

Syntax **rapid-retransmit-time** *hundred-milliseconds*
no rapid-retransmit-time

Context config>router>rsvp

Description This command defines the value of the Rapid Retransmission Interval. It is used in the re-transmission mechanism to handle unacknowledged message_id objects and is based on an exponential back-off timer.

Re-transmission interval of a RSVP message with the same message_id = 2 * rapid-retransmit-time interval of time.

The node stops re-transmission of unacknowledged RSVP messages:

- If the updated back-off interval exceeds the value of the regular refresh interval.
- If the number of re-transmissions 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

Parameters *hundred-milliseconds* — Specifies the rapid retransmission interval, in hundred-milliseconds (for example, enter “6” for a 600 millisecond retransmit time).

Values 1 – 100, in units of 100 ms.

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 message with the same message_id is retransmitted every 2 * rapid-retransmit-time interval of time. The node will stop retransmission of unacknowledged RSVP messages whenever the updated backoff interval exceeds the value of the regular refresh interval or the number of retransmissions reaches the value of the rapid-retry-limit parameter, whichever comes first.

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

Default 3

Parameters *number* — Specifies the value of the Rapid Retry Limit.

Values 1 – 6, integer values

refresh-time

Syntax **refresh-time** *seconds*
no refresh-time

Context config>router>rsvp

Description The **refresh-time** controls the interval (in s), between the successive Path and Resv refresh messages. RSVP declares the session down after it misses **keep-multiplier** *number* consecutive refresh messages.

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

Default 30 s

Parameters *seconds* — The refresh time in s.

Values 1 to 65535

te-threshold-update

Syntax [**no**] **te-threshold-update**

Context config>router>rsvp

Description This command is used to control threshold-based IGP TE updates. The **te-threshold-update** command must enable IGP TE update based only on bandwidth reservation thresholds per interface and must block IGP TE update on bandwidth changes for each reservation. Threshold levels can be defined using the **te-up-threshold** and **te-down-threshold** commands at the global RSVP or per-interface level.

The **no** form of this command should reset te-threshold-update to the default value and disable threshold based update.

Default no te-threshold-update

on-cac-failure

Syntax [no] on-cac-failure

Context config>router>rsvp>te-threshold-update

Description This command is used to enable a CAC failure-triggered IGP update.

The **no** form of this command should reset on-cac-failure to the default value and disable the CAC failure-triggered IGP update.

Default no on-cac-failure

update-timer

Syntax update-timer *seconds*
no update-timer

Context config>router>rsvp>te-threshold-update

Description This command is to control timer-based IGP TE updates. Timer-based IGP updates can be enabled by specifying a non-zero time value. Default value of update-timer is 0.

The **no** form of this command should reset update-timer to the default value and disable timer-based IGP update.

Default no update-timer (time - 0 s)

Parameters *seconds* — The time in s.

Values 0-300

te-up-threshold

Syntax te-up-threshold *threshold-level* [*threshold-level...*(up to 16 max)]
no te-up-threshold

Context config>router>rsvp
config>router>rsvp>interface

Description	This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The te-threshold-update command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates must be supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels must be supported. Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface is used to trigger IGP updates if non-default threshold levels are configured. The no form of this command resets te-up-threshold to its default value.
Default	0 15 30 45 60 75 80 85 90 95 96 97 98 99 100
Parameters	<i>threshold-level</i> — Integer value Values 0 to 100

te-down-threshold

Syntax	te-down-threshold <i>threshold-level</i> [<i>threshold-level</i> ...(up to 16 max)] no te-down-threshold
Context	config>router>rsvp config>router>rsvp>interface
Description	This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The te-threshold-update command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates must be supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels is supported. Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface is used to trigger IGP updates if non-default threshold levels are configured. The no form of this command resets te-down-threshold to its default value.
Default	100 99 98 97 96 95 90 85 80 75 60 45 30 15 0
Parameters	<i>threshold-level</i> — Integer value Values 0 to 100

2.20.2.2.3 Interface Commands

interface

Syntax	[no] interface <i>ip-int-name</i>
Context	config>router>rsvp
Description	This command enables RSVP protocol support on an IP interface. No RSVP commands are executed on an IP interface where RSVP is not enabled. The no form of this command deletes all RSVP commands such as hello-interval and subscription, which are defined for the interface. The RSVP interface must be shutdown it can be deleted. If the interface is not shut down, the no interface <i>ip-int-name</i> command does nothing except issue a warning message on the console indicating that the interface is administratively up.
Default	shutdown
Parameters	<i>ip-int-name</i> — The name of the network IP interface. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.
Values	1 to 32 alphanumeric characters.

authentication-key

Syntax	authentication-key [<i>authentication-key</i> <i>hash-key</i>] [hash hash2] no authentication-key
Context	config>router>rsvp>interface
Description	This command specifies the authentication key to be used between RSVP neighbors to authenticate RSVP messages. Authentication uses the MD-5 message-based digest. When enabled on an RSVP interface, authentication of RSVP messages operates in both directions of the interface. A node maintains a security association using one authentication key for each interface to a neighbor. The following items are stored in the context of this security association: • The HMAC-MD5 authentication algorithm. • Key used with the authentication algorithm. • Lifetime of the key. The user-entered key is valid until the user deletes it from the interface. • Source Address of the sending system. • Latest sending sequence number used with this key identifier.

A router RSVP sender transmits an authenticating digest of the RSVP message, computed using the shared authentication key and a keyed-hash algorithm. The message digest is included in an integrity object which also contains a flags field, a key identifier field, and a sequence number field. The RSVP sender complies to the procedures for RSVP message generation in RFC 2747, *RSVP Cryptographic Authentication*.

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

When a PLR node switches the path of the LSP to a bypass LSP, it does not send the Integrity object in the RSVP messages sent over the bypass tunnel. If the PLR receives an RSVP 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 message is a RESV message and does not have an Integrity object, then the PLR node will accept it only if it originated from the MP node.

An MP node will accept RSVP 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 MD5 implementation does not support the authentication challenge procedures in RFC 2747.

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

Default	no authentication-key - The authentication key value is the null string.
Parameters	<i>authentication-key</i> — The authentication key. The key can be any combination of ASCII characters up to 16 characters in length (unencrypted). If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. <i>hash-key</i> — 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 or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the hash2 encrypted variable cannot be copied and pasted. If the hash or hash2 parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash or hash2 parameter specified.

auth-keychain

Syntax	auth-keychain <i>name</i>
Context	config>router>rsvp>interface
Description	This command configures an authentication keychain to use for authentication of protocol messages sent and received over the associated interface. The keychain must include a valid entry to properly authenticate protocol messages, including a key, specification of a supported authentication algorithm, and beginning time. Each entry may also include additional options to control the overall lifetime of each entry to allow for the seamless rollover of without affecting the protocol adjacencies. The no form of the auth-keychain command removes the association between the routing protocol and any keychain currently used.
Default	no auth-keychain
Parameters	<i>name</i> — Specifies the name of the keychain, up to 32 characters, to use for the specified protocol session or sessions.

bfd-enable

Syntax	[no] bfd-enable
Context	config>router>rsvp>interface
Description	This command enables the use of bi-directional forwarding (BFD) to control the state of the associated RSVP interface. This causes RSVP to register the interface with the BFD session on that interface. The user configures the BFD session parameters, such as, transmit-interval, receive-interval, and multiplier, under the IP interface in the config>router> if>bfd context.



Note: It is possible that the BFD session on the interface was started because of a prior registration with another protocol, for example, OSPF or IS-IS.

The registration of an RSVP interface with BFD is performed at the time of neighbor gets its first session. This means when this node sends or receives a new Path message over the interface. If however the session did not come up, due to not receiving a Resv for a new path message sent after the maximum number of re-tries, the LSP is shutdown and the node de-registers with BFD. In general, the registration of RSVP with BFD is removed as soon as the last RSVP session is cleared.

The registration of an RSVP interface with BFD is performed independent of whether RSVP hello is enabled on the interface or not. However, hello timeout will clear all sessions towards the neighbor and RSVP de-registers with BFD at clearing of the last session.

An RSVP session is associated with a neighbor based on the interface address the path message is sent to. If multiple interfaces exist to the same node, each interface is treated as a separate RSVP neighbor. The user will have to enable BFD on each interface and RSVP will register with the BFD session running with each of those neighbors independently

Similarly the disabling of BFD on the interface results in removing registration of the interface with BFD.

When a BFD session transitions to DOWN state, the following actions are triggered. For RSVP signaled LSPs, this triggers activation of FRR bypass/detour backup (PLR role), global revertive (head-end role), and switchover to secondary if any (head-end role) for affected LSPs with FRR enabled. It triggers switchover to secondary if any and scheduling of re-tries for signaling the primary path of the non-FRR affected LSPs (head-end role).

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

Default no bfd-enable

hello-interval

Syntax **hello-interval** *milli-seconds*
no hello-interval

Context config>router>rsvp>interface

Description This command configures the time interval between RSVP hello messages.

RSVP hello packets are used to detect loss of RSVP connectivity with the neighboring node. Hello packets detect the loss of neighbor far quicker than it would take for the RSVP session to time out based on the refresh interval. After the loss of the of number keep-multiplier 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 ms

Parameters *milli-seconds* — Specifies the RSVP hello interval (in ms), in multiples of 1000. A 0 (zero) value disables the sending of RSVP hello messages.

Values 0 to 60000 ms (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 the use of the implicit null label over a specific RSVP interface. All LSPs for which this node is the egress LER and for which the path message is received from the previous hop node over this RSVP interface will signal the implicit null label. This means that if the egress LER is also the merge-point (MP) node, then the incoming interface for the path refresh message over the bypass dictates if the packet will use the implicit null label or not. The same for a 1-to-1 detour LSP. The user must shutdown the RSVP interface before being able to change the implicit null configuration option. The no form of this command returns the RSVP interface to use the RSVP level configuration value.
Default	disable
Parameters	<i>enable</i> — This parameter enables the implicit null label. <i>disable</i> — This parameter disables the implicit null label.

refresh-reduction

Syntax	[no] refresh-reduction
Context	config>router>rsvp>interface
Description	This command enables the use of the RSVP overhead refresh reduction capabilities on this RSVP interface. When this option is enabled, a node will enable support for three capabilities. It will accept bundles RSVP messages from its peer over this interface, it will attempt to perform reliable RSVP message delivery to its peer, and 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 message is intended to reduce 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 message except another bundle message. A node will only process received bundled RSVP messages but will not generate them. When reliable message delivery is supported by both the node and its peer over the RSVP interface, an RSVP message is sent with a message_id object. A message_id object can be added to any RSVP message when sent individually or as a sub-message of a bundled message. if the sender sets the ack_desired flag in the message_id object, the receiver acknowledges the receipt of the RSVP message by piggy-backing a message_ack object to the next RSVP 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 router supports the sending of separate ACK messages only but is capable of processing received message_ack objects piggy-backed to hop-by-hop RSVP messages, such as path and resv.

The router sets the ack_desired flag only in non refresh RSVP messages and in refresh messages which 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 message with the same message_id is retransmitted every 2 * rapid-retransmit-time interval of time. The rapid-retransmit-time is referred to as the rapid retransmission interval as 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 message rapid-retry-limit. The node will stop retransmission of unacknowledged RSVP messages whenever the updated backoff interval exceeds the value of the regular refresh interval or the number of retransmissions reaches the value of the rapid-retry-limit parameter, whichever comes first. These two parameters are configurable globally on a system in the **config>router>rsvp** context.

Refresh summary 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_id 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_id 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 an RSVP interface, the node indicates this to its peer by setting a “refresh-reduction-capable” bit in the flags field of the common RSVP header. If both peers of an RSVP interface set this bit, all the above three capabilities can be used. Furthermore, the node monitors the settings of this bit in received RSVP messages from the peer on the interface. As soon as this bit is cleared, the router 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, 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 an error message “unknown object class”. In this case, the node retransmits the RSVP message without the message_id object and reverts to using this method for future messages destined to this peer. The RSVP Overhead Refresh Reduction is supported with both RSVP P2P LSP path and the S2L path of an RSVP P2MP LSP instance over the same RSVP instance.

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 messages over the RSVP interface. When refresh-reduction is enabled on an interface and reliable-delivery is disabled, the router will send a message_id and not set ACK desired in the RSVP messages over the interface. The router does not expect an ACK and 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 header, the node will enter summary refresh for a specific message_id it sent regardless if it received an ACK or not to this message from the neighbor. Finally, when ‘reliable-delivery’ option is enabled on any interface, RSVP message pacing is disabled on all RSVP interfaces of the system, for example, the user cannot enable the msg-pacing option in the config>router>rsvp context, and error message is returned in CLI. Conversely, when the msg-pacing option is enabled, the user cannot enable the reliable delivery option on any interface on this system. An error message will also 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 <i>percentage</i> no subscription
Context	config>router>rsvp>interface
Description	This command configures the percentage of the link bandwidth that RSVP can use for reservation and sets a limit for the amount of over-subscription or under-subscription allowed on the interface. When the subscription is set to zero, no new sessions are permitted on this interface. If the <i>percentage</i> is exceeded, the reservation is rejected and a log message is generated. The no form of this command reverts the <i>percentage</i> to the default value.
Default	100
Parameters	<i>percentage</i> — The percentage of the interface's bandwidth that RSVP allows to be used for reservations. Values 0 to 1000

te-up-threshold

Syntax	te-up-threshold <i>threshold-level</i> [<i>threshold-level...</i> (up to 16 max)] no te-up-threshold
Context	config>router>rsvp config>router>rsvp>interface
Description	This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The te-threshold-update command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates must be supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels must be supported. Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface is used to trigger IGP updates if non-default threshold levels are configured. The no form of this command resets the default value.
Default	0 15 30 45 60 75 80 85 90 95 96 97 98 99 100
Parameters	<i>threshold-level</i> — Integer value
	Values 0 to 100

te-down-threshold

Syntax	te-down-threshold <i>threshold-level</i> [<i>threshold-level...</i> (up to 16 max)] no te-down-threshold
Context	config>router>rsvp config>router>rsvp>interface
Description	This command configures the specific threshold levels per node and per interface. Threshold levels are for reserved bandwidth per interface. The te-threshold-update command is used to enable or disable threshold-based IGP TE updates. Any reserved bandwidth change per interface is compared with all the threshold levels and trigger an IGP TE update if a defined threshold level is crossed in either direction (LSP setup or teardown). Threshold-based updates is supported with both ISIS and OSPF. A minimum of one and a maximum of 16 threshold levels is supported. Threshold levels configured per node is inherited by all configured RSVP interfaces. Threshold levels defined under the RSVP interface must be used to trigger IGP updates if non-default threshold levels are configured. The no form of this command resets the default value.

Default	100 99 98 97 96 95 90 85 80 75 60 45 30 15 0
Parameters	<i>threshold-level</i> — Integer value
Values	0 to 100

Message Pacing Commands

msg-pacing

Syntax	[no] msg-pacing
Context	config>router>rsvp
Description	This command enables RSVP message pacing in which the specified number of RSVP messages, specified in the max-burst command, are sent in a configured interval, specified in the period command. A count is kept of the messages that were dropped because the output queue for the interface used for message pacing was full.
Default	no msg-pacing

max-burst

Syntax	max-burst <i>number</i> no max-burst
Context	config>router>rsvp>msg-pacing
Description	This command specifies the maximum number of RSVP messages that are sent in the specified period under normal operating conditions.
Default	650
Parameters	<i>number</i> — The maximum number of RSVP messages to be sent
Values	100 to 1000 in increments of 10

period

Syntax	period <i>milli-seconds</i> no period
Context	config>router>rsvp>msg-pacing
Description	This command specifies the time interval (in ms), when the router can send the specified number of RSVP messages which is specified in the max-burst command.

Default	100
Parameters	<i>milli-seconds</i> — The time interval (in ms)
Values	100 to 1000 ms in increments of 10 ms

2.21 MPLS/RSVP Show, Tools, Router, Clear, and Debug Command Reference

2.21.1 Command Hierarchies

- [Show Commands](#)
- [Tools Commands](#)
- [Router Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)

2.21.1.1 Show Commands

```

show
  — router
    — mpls
      — bypass-tunnel [to ip-address] [protected-lsp name] [dynamic | manual |
        p2mp] [detail]
      — interface [ip-int-name | ip-address] [label-map label]
      — interface [ip-int-name | ip-address] statistics
      — label start-label [end-label | in-use | owner]
      — label-range
      — 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 [lsp-name] auto-bandwidth [auto-lsp {all | mesh-p2p | one-hop-p2p}]
      — lsp [lsp-name] path [path-name] mbb
      — lsp-egress-stats
      — lsp-egress-stats lsp-name
      — lsp-egress-stats ip-address lsp lsp-name
      — lsp-egress-stats sender-address:lsp-name
      — lsp-template [lsp-template-name] [detail]
      — mpls-tp
        — oam-template
        — protection-template
        — status

```

```

    — transit-path [path-name] [detail]
  — p2mp-info {type {originate | transit | terminate}} [s2l-endpoint ip-address]
  — p2mp-lsp [lsp-name] [detail]
  — p2mp-lsp [lsp-name] p2mp-instance [p2mp-instance-name] [mbb]
  — p2mp-lsp [lsp-name] p2mp-instance [p2mp-instance-name] s2l [s2l-name [to
    s2l-to-address]][status {up | down}] [detail]
  — p2mp-lsp [lsp-name] p2mp-instance [p2mp-instance-name] s2l [s2l-name [to
    s2l-to-address]] mbb
  — sr-te-lsp [lsp-name] [status {up | down}] [detail] path [path-name]
  — sr-te-lsp [lsp-name] [detail]
  — sr-te-lsp [lsp-name] [status {up | down}] [to ip-address] [detail]
  — srlg-database [router-id ip-address] [interface ip-address]
  — srlg-group [group-name]
  — static-lsp [lsp-name]
  — static-lsp {transit | terminate}
  — static-lsp count
  — statistics-summary
  — status
  — tp-lsp [lsp-name] [status {up | down}] [from ip-address | to ip-address] [detail]
  — tp-lsp [lsp-name] path [protecting | working] [detail]
  — tp-lsp [lsp-name] protection

show
  — router
    — bfd
      — session {ipv4 | ipv6} detail [lag lag-id] lag-port port-id
      — session lsp-name Lsp Name
      — session lsp-rsvp {head | tail}
      — session src ip-address/link-local address dest ip-address | link-local address
        detail lsp-rsvp {head | tail} tunnel-id tunnel-id lsp-id lsp-id
      — session mpls-tp
      — session lsp-name Lsp Name [link-type {cc-only | cc-cv}] detail
      — session p2mp-interface interface-name detail
      — session src ip-address/link-local address detail lsp-rsvp {head | tail} rsvp-
        session-name [256 chars max]
      — session [src ip-address/link-local address] [ipv4 | ipv6]
      — session src ip-address/link-local address dest ip-address | link-local address
      — session src ip-address/link-local address detail
      — session summary
      — session type type [ipv4 | ipv6]
    — rsvp
      — interface [interface [ip-int-name]] statistics [detail]
      — neighbor [ip-address] [detail]
      — session [session-type] [from ip-address | to ip-address | lsp-name name]
        [status {up | down}][detail]
      — statistics
      — status

show
  — test-oam
    — lsp-bfd

```

- **lsp-bfd local-bfd-discrim** *bfd-discriminator*
- **lsp-bfd lsp-name** *lsp-name*

2.21.1.2 Tools Commands

```

tools
  — dump
    — router
      — mpls
        — bypass-tunnel [lsp-name] plr
        — ftn
        — ilm
        — lspinfo
        — memory-usage
        — te-lspinfo [endpoint ip-address] [sender ip-address] [lspid lsp-id]
          [detail] [p2p | p2p-tid tunnel-id]
        — te-lspinfo [endpoint ip-address] [sender ip-address] [lspid lsp-id]
          [detail] [p2mp | p2mp-tid tunnel-id]{ [phops] [nhops] [s2l ip-
            address] } }
        — tp-interface interface-num
      — rsvp
        — psb
        — rsb
  — perform
    — router
      — mpls
        — adjust-autobandwidth [lsp lsp-name] [force [bandwidth mbps]]
        — cspf to ip-addr [from ip-addr] [bandwidth bandwidth] [include-bitmap
          bitmap] [exclude-bitmap bitmap] [hop-limit limit] [exclude-
          address excl-addr [excl-addr...(up to 8 max)]] [use-te-metric]
          [strict-srlg] [srlg-group grp-id...(up to 8 max)] [exclude-node excl-
          node-id [excl-node-id...(up to 8 max)]] [skip-interface interface-
          name] [ds-class-type class-type] [cspf-reqtype req-type] [least-
          fill-min-thd thd] [setup-priority val] [hold-priority val]
        — force-switch-path [lsp lsp-name] [path path-name]
        — [no] force-switch-path [lsp lsp-name]
        — resignal [lsp lsp-name] path path-name [delay minutes]
        — resignal {p2mp-lsp p2mp-lsp-name p2mp-instance p2mp-instance-
          name | p2mp-delay p2mp-minutes}
        — resignal-bypass [lsp bypass-lsp-name] [force] [delay minutes]
        — revert [lsp lsp-name]
        — sr-te-cspf to ip-address [path path-name] [from ip-address]
        — switch-path [lsp lsp-name] [path path-name]
        — trap-suppress number-of-traps time-interval
        — update-path [lsp lsp-name] path current-path-name new-path new-
          path-name]

```

2.21.1.3 Router Commands

```

config
  — router
    — [no] igmp
      — [no] tunnel-interface rsvp-p2mp lsp-name [sender ip-address]
    — [no] tunnel-interface rsvp-p2mp lsp-name [sender ip-address]

```

2.21.1.4 Clear Commands

```

clear
  — router
    — bfd
      — session src-ip ip-address dst-ip ip-address
      — session src-ip ip-address dst-ip ip-address lsp-rsvp {head | tail} tunnel-id
        [0..4294967295] lsp-id [0..4294967295]
      — session mpls-tp lsp-name path {working | protect}
      — session p2mp-interface interface-name
      — session src-ip ip-address lsp-rsvp {head | tail} rsvp-session-name [256
        chars max]
      — statistics src-ip ip-address dst-ip ip-address
      — statistics all
      — statistics src-ip ip-address dst-ip ip-address lsp-rsvp {head | tail} tunnel-id
        [0..4294967295] lsp-id [0..4294967295]
      — statistics mpls-tp lsp-name path {working | protect}
      — statistics p2mp-interface interface-name
      — statistics src-ip ip-address lsp-rsvp {head | tail} rsvp-session-name [256
        chars max]
    — mpls
      — interface [ip-int-name] [statistics]
      — lsp lsp-name
      — lsp-autobandwidth [lsp-name]
      — lsp-egress-stats
      — lsp-egress-stats lsp-name
      — lsp-ingress-stats
      — lsp-ingress-stats ip-address lsp lsp-name
      — lsp-ingress-stats sender-address:lsp-name
    — rsvp
      — interface
      — statistics

```

2.21.1.5 Debug Commands

```

debug
  — router
    — mpls [lsp lsp-name] [sender source-address] [endpoint endpoint-address] [tunnel-id
      tunnel-id] [lsp-id lsp-id] [interface ip-int-name]

```

```

— no mpls
  — [no] event
    — all [detail]
    — no all
    — frr [detail]
    — no frr
    — iom [detail]
    — no iom
    — lsp-setup [detail]
    — no lsp-setup
    — mbb [detail]
    — no mbb
    — misc [detail]
    — no misc
    — xc [detail]
    — no xc
— rsvp [lsp lsp-name] [sender source-address] [endpoint endpoint-address] [tunnel-id
tunnel-id] [lsp-id lsp-id] [interface ip-int-name]
— no rsvp
  — [no] event
    — all [detail]
    — no all
    — misc [detail]
    — no misc
    — nbr [detail]
    — no nbr
    — path [detail]
    — no path
    — resv [detail]
    — no resv
    — te-threshold-update
    — no te-threshold-update
  — [no] packet
    — all [detail]
    — no all
    — 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

```

2.21.2 Command Descriptions

2.21.2.1 Show Commands



Note: The command outputs in this chapter are examples only; actual displays may differ depending on supported functionality and user configuration.

auto-lsp

Syntax	auto-lsp [<i>lsp-name</i>] auto-bandwidth auto-lsp [<i>lsp-name</i>] [status { up down }] [detail] [to <i>ip-address</i>] auto-lsp [<i>lsp-name</i>] [status { up down }] { mesh-p2p one-hop-p2p } [detail] [to <i>ip-address</i>]
Context	show>router>mpls
Description	This command displays Auto-LSP information.
Parameters	<i>lsp-name</i> — Specifies the LSP name. Values 80 characters max up down — Specifies the state. mesh-p2p one-hop-p2p — Specifies the auto LSP type.

Output

Sample Output

```
*A:Dut-C# show router mpls auto-lsp

=====
MPLS Auto-LSP Template
=====
LSP Name                               Type           Fastfail      Admin   Oper
                               Config          State         State
-----
meshP2pLsp3-10.20.1.6-61441           MeshP2P        Yes           Up      Up
meshP2pLsp2-10.20.1.1-61442           MeshP2P        Yes           Up      Up
meshP2pLsp2-10.20.1.2-61443           MeshP2P        Yes           Up      Up
meshP2pLsp2-10.20.1.4-61444           MeshP2P        Yes           Up      Up
meshP2pLsp2-10.20.1.5-61445           MeshP2P        Yes           Up      Up
meshP2pLsp2-10.20.1.6-61446           MeshP2P        Yes           Up      Up
meshP2pLsp10-10.20.1.1-61447          MeshP2P        Yes           Up      Up
meshP2pLsp10-10.20.1.2-61448          MeshP2P        Yes           Up      Up
```

h2>bypass-tunnel

Syntax	bypass-tunnel [to <i>ip-address</i>] [protected-lsp [<i>lsp-name</i>]] [dynamic manual p2mp] [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 which serves as a backup for a set of LSPs. Such an LSP tunnel is called a bypass tunnel.
Parameters	<i>ip-address</i> — Specify the IP address of the egress router. <i>lsp-name</i> — Specify 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. <i>detail</i> — Displays detailed information. p2mp — Displays P2MP bypass tunnel information.
Output	MPLS Bypass Tunnel Output Fields

[Table 15](#) describes MPLS bypass tunnel output fields.

Table 15 MPLS Bypass Tunnel Output Fields

Label	Description
To	The system IP address of the egress router.
State	The LSP's administrative state.
Out I/F	Specifies the name of the network IP interface.
Out Label	Specifies the incoming MPLS label on which to match.
Reserved BW (Kbps)	Specifies the amount of bandwidth in Mb/s reserved for the LSP.

h3>Sample Output

```
*A:Dut-B# show router mpls bypass-tunnel detail

=====
MPLS Bypass Tunnels (Detail)
=====
-----
bypass-node10.20.1.4
-----
To           : 10.20.1.7           State          : Up
Out I/F      : 1/1/4              Out Label      : 131071
Up Time     : 0d 01:17:22         Active Time    : n/a
Reserved BW  : 0 Kbps             Protected LSP Count : 1
Type        : Dynamic
```

```

Setup Priority : 7                      Hold Priority : 0
Class Type    : 0
Exclude Node   : 10.20.1.4             Inter-Area : True
Computed Hops  :
    10.10.8.2(S)                       Egress Admin Groups : None
-> 10.10.8.6(SA)                       Egress Admin Groups : None
-> 10.20.1.7(L)                         Egress Admin Groups : None
Actual Hops    :
    10.10.8.2(10.20.1.2)               Record Label : N/A
-> 10.10.8.6(10.20.1.6)               Record Label : 131071
-> 10.20.1.7(10.20.1.7)               Record Label : 131068
-> 10.10.22.7                         Record Label : 131068

```

```

=====
*A:Dut-A>config>router>mpls>lsp$ /show router mpls bypass-tunnel detail

```

```

=====
MPLS Bypass Tunnels (Detail)
=====
-----

```

```

bypass-node10.20.1.2
-----

```

```

To           : 10.20.1.4                State          : Up
Out I/F       : 1/1/2                  Out Label       : 131070
Up Time       : 0d 00:00:18            Active Time     : n/a
Reserved BW   : 0 Kbps                 Protected LSP Count : 1
Type          : Dynamic
Setup Priority : 7                      Hold Priority    : 0
Class Type    : 0
Exclude Node   : None                  Inter-Area      : False
Computed Hops  :
    10.20.1.1, If Index : 3(S)         Egress Admin Groups : None
-> 10.20.1.3, If Index : 2(S)         Egress Admin Groups : None
-> 10.20.1.4, If Index : 5(S)         Egress Admin Groups : None
Actual Hops    :
    10.20.1.1, If Index : 3            Record Label     : N/A
-> 10.20.1.3, If Index : 2            Record Label     : 131070
-> 10.20.1.4, If Index : 5            Record Label     : 131070

```

```

=====
B:Dut-B>config>router>mpls>lsp# show router mpls bypass-tunnel detail

```

```

=====
MPLS Bypass Tunnels (Detail)
=====
-----

```

```

bypass-node10.20.1.4
-----

```

```

To           : 10.10.10.6                State          : Up
Out I/F       : lag-1                  Out Label       : 131071
Up Time       : 0d 00:00:06            Active Time     : n/a
Reserved BW   : 0 Kbps                 Protected LSP Count : 1
Type          : Dynamic
Setup Priority : 7                      Hold Priority    : 0
Class Type    : 0
Exclude Node   : None

```



```

Actual Hops      :
    10.10.12.2 (S)

-> 10.10.12.3 (S)

-> 10.10.5.5 (S)

-> 10.10.10.6 (S)
  
```

Egress Admin Groups:

- lime
- olive
- blue
- black
- acqua

Egress Admin Groups:

- olive
- Unknown Group 9
- Unknown Group 11
- black
- Unknown Group 16
- Unknown Group 18

Egress Admin Groups:

- purple
- Unknown Group 7
- Unknown Group 11
- orange
- acqua
- Unknown Group 16
- Unknown Group 19
- Unknown Group 21
- Unknown Group 22
- Unknown Group 26
- khaki

Egress Admin Groups: None

=====:

*A:SRU4>show>router>mpls# bypass-tunnel

=====

MPLS Bypass Tunnels

Legend : m - Manual d - Dynamic p - P2mp

To	State	Out I/F	Out Label	Reserved BW (Kbps)	Protected LSP Count	Type

No Matching Entries Found

=====

*A:SRU4>show>router>mpls#

*A:Dut-B# show router mpls bypass-tunnel detail

=====

MPLS Bypass Tunnels (Detail)

=====

bypass-link10.10.104.4

To	: 10.10.101.4	State	: Up
Out I/F	: 1/1/2:1	Out Label	: 129994
Up Time	: 0d 00:02:33	Active Time	: n/a
Reserved BW	: 0 Kbps	Protected LSP Count	: 1
Type	: Dynamic		
SetupPriority	: 7	Hold Priority	: 0
Class Type	: 0		
Actual Hops	:		

```

10.10.101.2      -> 10.10.101.4
=====
*A:Dut-B#

*A:Dut-B# show router mpls bypass-tunnel detail
=====
MPLS Bypass Tunnels (Detail)
-----
bypass-link10.10.104.4
-----
To           : 10.10.101.4      State           : Up
Out I/F      : 1/1/2:1         Out Label      : 129994
Up Time      : 0d 00:02:33     Active Time    : n/a
Reserved BW  : 0 Kbps         Protected LSP Count : 1
Type         : Dynamic
SetupPriority : 7              Hold Priority   : 0
Class Type   : 0
Actual Hops   :
10.10.101.2      -> 10.10.101.4
=====
*A:Dut-B#

```

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* — The name of the network IP interface. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.
- ip-address* — The system or network interface IP address.
- label-map label* — The MPLS label on which to match.
- Values** 32 to 1048575
- statistics* — Displays the MPLS interface name and the number of packets and octets sent and received on an MPLS interface.
- Output** MPLS Interface Output Fields
- [Table 16](#) describes MPLS interface output fields.

Table 16 MPLS Interface Output Fields

Label	Description
Interface	The interface name.

Table 16 MPLS Interface Output Fields (Continued)

Label	Description (Continued)
Port-id	The port ID.
Adm	Specifies the administrative state of the interface.
Opr	Specifies the operational state of the interface.
Te-metric	Specifies the traffic engineering metric used on the interface.
Srlg Groups	Specifies the shared risk loss group (SRLG) name(s).
Interfaces	The total number of interfaces.
Transmitted	Displays the number of packets and octets transmitted from the interface.
Received	Displays the number of packets and octets received.
In Label	Specifies the ingress label.
In I/F	Specifies the ingress interface.
Out Label	Specifies the egress label.
Out I/F	Specifies the egress interface.
Next Hop	Specifies the next hop IP address for the static LSP.
Type	Specifies whether the label value is statically or dynamically assigned.

Sample Output

```
*A:SRU4>config>router>mpls# show router mpls interface
=====
MPLS Interfaces
=====
Interface                Port-id                Adm    Opr    TE-metric
-----
system                   system                 Up     Up     None
  Admin Groups           None
  Srlg Groups            None
sr4-1                    1/1/4                 Up     Up     None
  Admin Groups           None
  Srlg Groups            3440
ess-7-1                  3/2/4                 Up     Up     None
  Admin Groups           None
  Srlg Groups            45100
ess-7-2                  3/2/5                 Up     Up     None
  Admin Groups           None
  Srlg Groups            45110
...
g7600                    3/1/2                 Up     Up     None
  Admin Groups           None
  Srlg Groups            41.80
```

```

m160                                3/2/1                Up    Up    None
  Admin Groups                     None
  Srlg Groups                       420.40
-----
Interfaces : 35
=====
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls interface "hubA"
=====
MPLS Interface : hubA
=====
Interface                Port-id            Adm    Opr    TE-metric
-----
hubA                      3/2/8              Up     Up     None
  Admin Groups            None
  Srlg Groups              44.200
-----
Interfaces : 1
=====
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls interface "hubA" label-map 203
=====
MPLS Interface : hubA (Label-Map 203)
=====
In Label  In I/F      Out Label Out I/F      Next Hop      Type      Adm  Opr
-----
203        3/2/8        403        1/1/9        11.22.10.3    Static    Up   Up
-----
Interfaces : 1
=====
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls interface statistics
=====
MPLS Interface (statistics)
=====
Interface      : aps-1
  Transmitted   : Pkts - 76554                Octets - 7930285
  Received      : Pkts - 17068              Octets - 3626842

Interface      : aps-2
  Transmitted   : Pkts - 0                  Octets - 0
  Received      : Pkts - 1311              Octets - 219888

Interface      : aps-3
  Transmitted   : Pkts - 0                  Octets - 0
  Received      : Pkts - 3                 Octets - 234

Interface      : sr4-1
  Transmitted   : Pkts - 0                  Octets - 0
  Received      : Pkts - 0                  Octets - 0

Interface      : ess-7-1
  Transmitted   : Pkts - 113537             Octets - 15058332

```

```

Received      : Pkts - 13193                Octets - 1091492

Interface     : ess-7-2
  Transmitted  : Pkts - 166133                Octets - 22762482
  Received    : Pkts - 16672                Octets - 1368464

Interface     : ess-7-3
  Transmitted  : Pkts - 122934                Octets - 11033246
  Received    : Pkts - 12256                Octets - 1026826
...

Interface     : m160
  Transmitted  : Pkts - 17188024            Octets - 2183076528
  Received    : Pkts - 677745              Octets - 59367236
=====
*A:SRU4>config>router>mpls#

```

```

*A:SRU4>config>router>mpls# show router mpls interface

```

```

=====
MPLS Interfaces
=====

```

Interface	Port-id	Adm	Opr	TE-metric
system	system	Up	Up	None
Admin Groups	None			
Srlg Groups	None			
aps-1	aps-1	Up	Up	None
Admin Groups	None			
Srlg Groups	3410			
aps-2	aps-2	Up	Up	None
Admin Groups	None			
Srlg Groups	3420			
aps-3	aps-3	Up	Up	None
Admin Groups	None			
Srlg Groups	3430			
sr4-1	1/1/4	Up	Up	None
Admin Groups	None			
Srlg Groups	3440			
ess-7-1	3/2/4	Up	Up	None
Admin Groups	None			
Srlg Groups	45100			
ess-7-2	3/2/5	Up	Up	None
Admin Groups	None			
Srlg Groups	45110			
...				
g7600	3/1/2	Up	Up	None
Admin Groups	None			
Srlg Groups	41.80			
m160	3/2/1	Up	Up	None
Admin Groups	None			
Srlg Groups	420.40			

```

-----
Interfaces : 35
=====

```

label

- Syntax** `label start-label [end-label | in-use | owner]`
- Context** `show>router>mpls`
- Description** Displays MPLS labels exchanged.
- Parameters**
 - start-label* — The label value assigned at the ingress router.
 - end-label* — The label value assigned for the egress router.
 - in-use* — The number of in-use labels displayed.

Output MPLS Label Output Fields

[Table 17](#) describes MPLS label output fields.

Table 17 MPLS Label Output Fields

Label	Description
Label	Displays the value of the label being displayed.
Label Type	Specifies whether the label value is statically or dynamically assigned.
Label Owner	The label owner.
In-use labels in entire range	The total number of labels being used by RSVP.

Sample Output

```
*A:mlstp-dutA# show router mpls label-range

=====
Label Ranges
=====
Label Type      Start Label      End Label      Aging      Total Available
-----
Static-lsp      32               16415         -          16364
Static-svc      16416            32799         -          16376
Dynamic         32800            131071        0          98268
=====

*A:SRU4>config>router>mpls# show router mpls label 202
=====
MPLS Label 202
=====
Label      Label Type      Label Owner
-----
202        static-lsp      STATIC
```

```
-----  
In-use labels in entire range          : 5057  
-----  
*A:SRU4>config>router>mpls#
```

label-range

- Syntax** **label-range**
- Context** show>router>mpls
- Description** This command displays the MPLS label range.
- Output** MPLS Label Range Output

[Table 18](#) describes the MPLS label range output fields.

Table 18 MPLS Label Range Output Fields

Label	Description
Label Type	Displays the 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.
Aging	The number of labels released from a service which are transitioning back to the label pool. Labels are aged 15 s.
Total Available	The number of label values available.

Sample Output

```
*A:SRU4>config>router>mpls# show router mpls label-range  
=====
```

Label Type	Start Label	End Label	Aging	Total Available
Static-lsp	32	1023	-	736
Static-svc	2048	18431	-	16384
Dynamic	32768	131071	258	93232

```
=====
```

*A:SRU4>config>router>mpls#

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** 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 [*lsp-name*] **auto-bandwidth** [**auto-lsp** {**all** | **mesh-p2p** | **one-hop-p2p**}]
lsp [*lsp-name*] **path** [*path-name*] **mbb**

Context show>router>mpls

Description This command displays LSP details.

Parameters *lsp-name* — 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 number of static LSPs that transit through the router.
terminate — Displays the number of static LSPs that terminate at the router.
lsp count — Displays the total number of LSPs.
activepath — Displays the present path being used to forward traffic.
mbb — Displays make-before-break (MBB) information.
detail — Displays detailed information.

Output MPLS LSP Output

[Table 19](#) describes MPLS LSP output fields.

Table 19 MPLS LSP Output Fields

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.
Adm State	Down — The path is administratively disabled. Up — The path is administratively enabled.

Table 19 MPLS LSP Output Fields (Continued)

Label	Description (Continued)
Oper State	Down — The path is operationally down. Up — The path is operationally up.
LSPs	The total number of LSPs configured.
From	The IP address of the ingress router for the LSP.
LSP Up Time	The length of time the LSP has been operational.
Transitions	The number of transitions that have occurred for the LSP.
Retry Limit	The number of attempts that the software should make to re-establish the LSP after it has failed.
Signaling	Specifies the signaling style.
Hop Limit	The maximum number of hops that an LSP can traverse, including the ingress and egress routers.
Fast Reroute/ FastFail Config	enabled — Fast reroute is enabled. In the event of a failure, traffic is immediately rerouted on the pre-computed detour LSP, thus minimizing packet loss. disabled — There is no detour LSP from each node on the primary path.
ADSPEC	enabled — The LSP will include advertising data (ADSPEC) objects in RSVP messages. disabled — The LSP will not include advertising data (ADSPEC) objects in RSVP messages.
Primary	The preferred path for the LSP.
Secondary	The alternate path that the LSP will use if the primary path is not available.
Bandwidth	The amount of bandwidth in Mb/s reserved for the LSP path.
LSP Up Time	The total time in increments that the LSP path has been operational.
LSP Tunnel ID	The value which identifies the label switched path that is signaled for this entry.
To	The IP address of the egress router for the LSP.
LSP Down Time	The total time in increments that the LSP path has not been operational.
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 Timer	The time in s, for LSP re-establishment attempts after an LSP failure.
Resv Style	se — Specifies a shared reservation environment with a limited reservation scope. This reservation style creates a single reservation over a link that is shared by an explicit list of senders. ff — Specifies a shared reservation environment with an explicit reservation scope. Specifies an explicit list of senders and a distinct reservation for each of them.

Table 19 MPLS LSP Output Fields (Continued)

Label	Description (Continued)
Negotiated MTU	The size of the maximum transmission unit (MTU) that is negotiated during establishment of the LSP.
FR Hop Limit	The total number of hops a detour LSP can take before merging back onto the main LSP path.
LastResignalAttempt	Displays the system up time when the last attempt to resignal this LSP was made.
MBB Type	Displays an enumerated integer that specifies the type of make-before-break (MBB). If none displays then there is no MBB in progress or no last MBB.
MBB State	Displays the state of the most recent invocation of the make-before-break functionality.
End at	Displays the system up time when the last MBB ended.
Old Metric	Displays the cost of the traffic engineered path for the LSP path prior to MBB.
NextRetryIn	Displays the amount of time (in s) remaining, before the next attempt is made to retry the in-progress MBB.
RetryAttempt	Displays the number attempts for the MBB is in progress.
Failure Code	Displays the reason code for in-progress MBB failure. A value of none indicates that no failure has occurred.
Failure Node	Displays 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 .

Sample Output

```

*A:SetupCLI# show router mpls lsp "lsp_1" path "500" detail
=====
MPLS LSP lsp_1 Path 500 (Detail)
=====
Legend :
    @ - Detour Available          # - Detour In Use
    b - Bandwidth Protected       n - Node Protected
    s - Soft Preemption
    S - Strict                    L - Loose
    A - ABR
=====
-----
LSP lsp_1 Path 500
-----
LSP Name       : lsp_1
Path LSP ID    : 38400
From           : 10.10.1.1      To           : 10.10.1.2
Admin State    : Up             Oper State    : Down
Path Name      : 500            Path Type    : Primary
Path Admin     : Up             Path Oper    : Down
Out Interface  : n/a            Out Label    : n/a

```

Path Up Time	: 0d 00:00:00	Path Down Time	: 0d 00:00:43
Retry Limit	: 0	Retry Timer	: 30 sec
Retry Attempt	: 2	Next Retry In	: 19 sec
BFD Template	: None	BFD Ping Interval	: 60
BFD Enable	: FALSE		
Adspec	: Disabled	Oper Adspec	: N/A
CSPF	: Disabled	Oper CSPF	: N/A
Least Fill	: Disabled	Oper LeastFill	: N/A
FRR	: Disabled	Oper FRR	: N/A
Prop Adm Grp	: Disabled	Oper Prop AG	: N/A
Inter-area	: N/A		
Neg MTU	: 0	Oper MTU	: N/A
Bandwidth	: No Reservation	Oper Bw	: N/A
Hop Limit	: 255	Oper HopLimit	: N/A
Record Route	: Record	Oper Rec Route	: N/A
Record Label	: No Record	Oper Rec Label	: N/A
Setup Priority	: 7	Oper Setup Priority	: N/A
Hold Priority	: 7	Oper Hold Priority	: N/A
Class Type	: 7	Oper CT	: N/A
Backup CT	: 5		
MainCT Retry	: 500		
Rem	:		
MainCT Retry	: 500		
Limit	:		
Include Grps	:	Oper InclGrps	:
None	:	N/A	:
Exclude Grps	:	Oper ExclGrps	:
None	:	N/A	:
Adaptive	: Enabled	Oper Metric	: N/A
Preference	: n/a		
Path Trans	: 0	CSPF Queries	: 0
Failure Code	: noResourcesAvailable		
Failure Node	: 9.1.255.255		
Explicit Hops	:		
No Hops Specified	:		
Actual Hops	:		
No Hops Specified	:		
Resignal Eligible	: False		
Last Resignal	: n/a	CSPF Metric	: N/A

=====

*A:Dut-A>config>router>mpls>lsp\$ /show router mpls lsp "1" path detail

=====

MPLS LSP 1 Path (Detail)

=====

Legend :

@ - Detour Available	# - Detour In Use
b - Bandwidth Protected	n - Node Protected
s - Soft Preemption	
S - Strict	L - Loose
A - ABR	

=====

LSP 1 Path 1

```

LSP Name      : 1
From          : 10.20.1.1
Adm State     : Up
Path Name     : 1
Path Admin    : Up
OutInterface  : 1/1/1
Path Up Time  : 0d 00:00:05
Retry Limit   : 0
RetryAttempt  : 0

Path LSP ID   : 30208
To            : 10.20.1.6
Oper State    : Up
Path Type     : Primary
Path Oper     : Up
Out Label     : 131071
Path Dn Time  : 0d 00:00:00
Retry Timer   : 30 sec
NextRetryIn   : 0 sec

Adspec        : Disabled
CSPF          : Enabled
Least Fill    : Disabled
FRR           : Enabled
FRR NodePro*  : Enabled
FR Hop Limit  : 16
FR Prop Adm*  : Disabled
Prop Adm Grp  : Disabled
Inter-area    : False

Oper Adspec   : Disabled
Oper CSPF     : Enabled
Oper LeastF*  : Disabled
Oper FRR      : Enabled
Oper FRR NP   : Enabled
Oper FRHopL*  : 16
Oper FRProp*  : Disabled
Oper PropAG   : Disabled

Neg MTU       : 1496
Bandwidth     : No Reservation
Hop Limit     : 255
Record Route  : Record
Record Label  : Record
SetupPriori*  : 7
Hold Priori*  : 0
Class Type    : 0
Backup CT     : None
MainCT Retry  : n/a
Rem           :
MainCT Retry  : 0
Limit         :
Include Grps  :
None
Exclude Grps :
None

Oper MTU       : 1496
Oper Bw        : 0 Mbps
Oper HopLim*   : 255
Oper RecRou*   : Record
Oper RecLab*   : Record
Oper SetupP*   : 7
Oper HoldPr*   : 0
Oper CT        : 0

Oper InclGr*   :
None
Oper ExclGr*   :
None

Adaptive      : Enabled
Preference    : n/a
Path Trans    : 1
Failure Code  : noError
ExplicitHops  :
    No Hops Specified
Actual Hops   :
    10.20.1.1, If Index : 2 @ n
-> 10.20.1.2, If Index : 2 @ n
-> 10.20.1.4, If Index : 2
-> 10.20.1.6, If Index : 2
ComputedHops  :
    10.20.1.1, If Index : 2(S)
-> 10.20.1.2, If Index : 2(S)
-> 10.20.1.4, If Index : 2(S)
-> 10.20.1.6, If Index : 2(S)
ResigEligib*  : False
LastResignal  : n/a

Oper Metric    : 3000

CSPF Queries  : 1
Failure Node   : n/a

Record Label   : N/A
Record Label   : 131071
Record Label   : 131071
Record Label   : 131071

CSPF Metric    : 3000
=====
* indicates that the corresponding row element may have been truncated.

```

```
*A:Dut-C# show router mpls lsp detail
=====
MPLS LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name      : to_D_10.20.1.4_viaBD
LSP Type      : RegularLsp
From          : 10.20.1.3
Adm State     : Up
LSP Up Time   : 0d 00:05:38
Transitions   : 1
Retry Limit   : 0
Signaling     : RSVP
Hop Limit     : 255
Adaptive      : Enabled
FastReroute   : Disabled
CSPF          : Enabled
Metric        : 0
Include Grps  :
None
Least Fill    : Disabled

LSP Tunnel ID : 1
To             : 10.20.1.4
Oper State     : Up
LSP Down Time : 0d 00:00:00
Path Changes   : 1
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : 1500
ClassType      : 0
Oper FR        : Disabled
ADSPEC         : Disabled
Use TE metric  : Disabled
Exclude Grps   :
None

Auto BW        : Disabled
LdpOverRsvp    : Enabled
IGP Shortcut   : Enabled
IGP LFA        : Disabled
BGPTransTun    : Enabled
Oper Metric    : 20
Prop Adm Grp   : Disabled

VprnAutoBind   : Enabled
BGP Shortcut   : Enabled
IGP Rel Metric : -1

Primary(a)     : to_D_10.20.1.4_viaBD
Bandwidth      : 0 Mbps
Up Time        : 0d 00:05:38
=====
```

*A:Dut-C#

```
*A:Dut-C# show router mpls lsp "lspE" detail
=====
MPLS LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name      : lspE
LSP Type      : RegularLsp
LSP Index     : 1005
From          : 10.20.1.3
Adm State     : Up
LSP Up Time   : 0d 00:19:23
Transitions   : 1
Retry Limit   : 0
Signaling     : RSVP
Hop Limit     : 255
Adaptive      : Enabled
FastReroute   : Disabled
CSPF          : Disabled
Metric        : N/A
Load Bal Wt   : N/A

LSP Tunnel ID : 1005
TTM Tunnel Id  : 1005
To             : 10.20.1.5
Oper State     : Up
LSP Down Time : 0d 00:00:00
Path Changes   : 1
Retry Timer    : 20 sec
Resv. Style    : SE
Negotiated MTU : 1500
ClassType      : 0
Oper FR        : Disabled
ADSPEC         : Disabled
ClassForwarding : Disabled
```

```

Include Grps:
None
Least Fill : Disabled
BFD Template: None
BFD Enable : False

Revert Timer: Disabled
EntropyLbl : Inherited
Auto BW : Disabled
LdpOverRsvp : Enabled
VprnAutoBind: Enabled
IGP Shortcut: Enabled
IGP LFA : Disabled
BGPTransTun : Enabled
Oper Metric : 16777215
Prop Adm Grp: Disabled

Exclude Grps :
None
BFD Ping Intvl : 60

Next Revert In : N/A
Oper EntropyLbl: Enabled

BGP Shortcut : Enabled
IGP Rel Metric : Disabled

Primary(a) : E1
Bandwidth : 0 Mbps
Standby : E2
Bandwidth : 0 Mbps

Up Time : 0d 00:19:32
Down Time : 0d 00:19:46
=====

```

A:sim1>config>router>mpls>lsp\$ 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

```

=====

LSP 11 Path 1

```

LSP Name      : 11
From          : 10.20.1.1
Adm State     : Up
Path Name     : 1
Primary
Path Admin    : Up
OutInterface  : n/a
Path Up Time  : 0d 00:00:00
Retry Limit   : 0
RetryAttempt  : 0
SetupPriorit : 7
Preference    : n/a
Bandwidth     : No Reservation
Hop Limit     : 255
Backup CT     : None
MainCT Retry  : n/a
Rem           :
Oper CT       : None
Record Route  : Record
Oper MTU      : 0
Adaptive      : Enabled
Include Grps  :

Path LSP ID   : 30208
To            : 10.20.1.3
Oper State    : Down
Path Type     :
Path Oper     : Down
Out Label     : n/a
Path Dn Time  : 0d 00:00:02
Retry Timer   : 30 sec
NextRetryIn   : 7 sec (Fast)
Hold Priorit  : 0

Oper Bw       : 0 Mbps
Class Type    : 0

MainCT Retry  : 0
Limit         :

Record Label   : Record
Neg MTU        : 0
Oper Metric    : 65535
Exclude Grps   :

```

```
None
Path Trans : 2
Failure Code: noError
ExplicitHops:
  10.20.1.2(S)
Actual Hops :
  No Hops Specified
ResigEligib*: False
LastResignal: n/a
CSPF Queries: 0
Failure Node: n/a
CSPF Metric : 0
=====
```

```
*A:SRU4>config>router>mpls# show router mpls lsp path
=====
MPLS LSP Path (Detail)
=====
Legend :
  @ - Detour Available          # - Detour In Use
  b - Bandwidth Protected       n - Node Protected
  s - Soft Preemption
  S - Strict                    L - loose
=====
ExplicitHops:
  10.20.1.3(L)      -> 10.20.1.4(S)
Actual Hops :
  10.10.1.1(10.20.1.1)      Record Label : N/A
-> 10.10.1.2(10.20.1.2)      Record Label : 131071
-> 10.10.5.3(10.20.1.3)      Record Label : 131071
-> 10.10.7.4(10.20.1.4)      Record Label : 131071
-> 10.10.8.5(10.20.1.5)      Record Label : 131071
ComputedHops:
  10.10.1.1(S)      -> 10.10.1.2(S)      -> 10.10.5.3(S)
-> 10.20.1.4(S)      -> 10.20.1.5(L)
=====
```

```
*A:SRU4>config>router>mpls# show router mpls lsp
=====
MPLS LSPs (Originating)
=====
LSP Name                                To                                Fastfail    Adm    Opr
                                Config
-----
to_110_20_1_1_cspf                    110.20.1.1                        No          Up     Up
to_110_20_1_2_cspf                    110.20.1.2                        No          Up     Dwn
to_110_20_1_3_cspf                    110.20.1.3                        No          Up     Up
to_110_20_1_4_cspf                    110.20.1.4                        No          Up     Dwn
to_110_20_1_5_cspf                    110.20.1.5                        No          Up     Up
to_110_20_1_6_cspf                    110.20.1.6                        No          Up     Dwn
to_110_20_1_110_cspf                  110.20.1.110                     No          Up     Up
to_10_8_100_15_cspf                   10.8.100.15                      No          Up     Dwn
to_10_20_1_20_cspf                    10.20.1.20                       No          Up     Up
to_10_20_1_22_cspf                    10.20.1.22                       No          Up     Up
to_10_100_1_1_cspf                    10.100.1.1                       No          Up     Dwn
to_110_20_1_1_cspf_2                  110.20.1.1                       No          Up     Up
to_110_20_1_1_cspf_3                  110.20.1.1                       No          Up     Up
to_110_20_1_1_cspf_4                  110.20.1.1                       No          Up     Up
to_110_20_1_1_cspf_5                  110.20.1.1                       No          Up     Up
```

```

to_110_20_1_1_cspf_6      110.20.1.1      No      Up      Up
to_110_20_1_1_cspf_7      110.20.1.1      No      Up      Up
to_110_20_1_1_cspf_8      110.20.1.1      No      Up      Up
to_110_20_1_1_cspf_9      110.20.1.1      No      Up      Up
to_110_20_1_1_cspf_10     110.20.1.1      No      Up      Up
to_110_20_1_1_cspf_11     110.20.1.1      No      Up      Up
to_110_20_1_1_cspf_12     110.20.1.1      No      Up      Up
to_110_20_1_1_cspf_13     110.20.1.1      No      Up      Up
to_110_20_1_1_cspf_14     110.20.1.1      No      Up      Up
to_110_20_1_1_cspf_15     110.20.1.1      No      Up      Up
...
-----
LSPs : 201
=====
*A:SRU4>config>router>mpls#

```

Table 20 LSP Detail Output Fields

Label	Description
Auto BW	Enabled — Auto-bandwidth adjustment is configured on this LSP.
AB OpState	Up – Auto-bandwidth is operationally enabled on this LSP Down – Auto-bandwidth is operationally disabled on this LSP
Auto BW Min	The minimum bandwidth of the LSP that auto-bandwidth can request (in Mb/s).
Auto BW Max	The maximum bandwidth of the LSP that auto-bandwidth can request (in Mb/s).
AB Up Thresh	The percent threshold for increasing LSP bandwidth.
AB Down Thresh	The percent threshold for decreasing LSP bandwidth.
AB Up BW	The absolute bandwidth threshold for increasing LSP bandwidth (in Mb/s).
AB Down BW	The absolute bandwidth threshold for decreasing LSP bandwidth (in Mb/s).
AB Coll Intv	The auto-bandwidth collection interval.
AB Adj Mul	The adjust-multiplier for this LSP (may be configured or inherited).
AB Samp Mul	The sample-multiplier for this LSP (may be configured or inherited).
AB Adj Time	The adjust-multiplier times the collection-interval (in min).
AB Sample Time	The sample-multiplier times the collection-interval (in min).
AB Adj Cnt	The adjust count (number of whole collection intervals since the start of the current adjust interval).

Table 20 LSP Detail Output Fields (Continued)

Label	Description (Continued)
AB Samp Cnt	The sample count (number of whole collection intervals since the start of the current sample interval).
AB Last Adj	The system time of the last auto-bandwidth adjustment.
AB Next Adj	The approximate remaining time in the current adjust interval (adjust-multiplier – adjust count) times the collection interval (in min). This overstates the actual remaining time because the elapsed time in the current collection interval is not accounted for.
AB Adj Cause	The cause of the last auto-bandwidth adjustment: • none – no adjustment has occurred • manual • adj-count • overflow
AB Max AvgR*	The maximum average data rate in any sample interval of the current adjust interval.
AB Lst AvgR*	The average data rate measured in the sample interval that ended most recently.
AB Ovfl Lmt	The configured value of the auto-bandwidth overflow-limit.
AB Ovfl Cnt	The number of overflow samples since the last reset.
ABOvflThres	The percent threshold for declaring an overflow sample.
AB Ovfl BW	The absolute bandwidth threshold for declaring an overflow sample (in Mb/s).
AB Monitor BW	True – monitor bandwidth is enabled on the LSP. False – monitor bandwidth is not enabled on the LSP.

```

*A:SRU4>config>router>mpls# show router mpls lsp "to_110_20_1_1_cspf"
=====
MPLS LSPs (Originating)
=====
LSP Name                               To                               Fastfail   Adm   Opr
                                   Config
-----
to_110_20_1_1_cspf                     110.20.1.1                       No         Up    Up
-----
LSPs : 1
=====
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls lsp "to_110_20_1_1_cspf" detail
=====

```

```

MPLS LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name      : to_110_20_1_1_cspf
LSP Type      : RegularLsp
From          : 110.20.1.4
Adm State     : Up
LSP Up Time   : 0d 01:47:02
Transitions   : 11
Retry Limit   : 0
Signaling     : RSVP
Hop Limit     : 255
Adaptive      : Enabled
FastReroute   : Disabled
CSPF          : Enabled
Metric        : 0
Include Grps  :
None
Least Fill    : Disabled
LdpOverRsvp   : Enabled
IGP Shortcut  : Enabled
Oper Metric   : 1001

LSP Tunnel ID : 1
Oper State     : Up
LSP Down Time  : 0d 00:00:00
Path Changes   : 11
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : 1500
ClassType      : 0
Oper FR        : Disabled
ADSPEC         : Disabled
Use TE metric  : Disabled
Exclude Grps   :
None

VprnAutoBind   : Enabled

Primary(a)     : to_110_20_1_1
Bandwidth      : 0 Mbps
Up Time        : 0d 01:47:02
=====
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls lsp detail to 110.20.1.2
=====
MPLS LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name      : 1
LSP Type      : RegularLsp
From          : 0.0.0.0
Adm State     : Down
LSP Up Time   : 0d 00:00:00
Transitions   : 0
Retry Limit   : 0
Signaling     : RSVP
Hop Limit     : 255
Adaptive      : Enabled
FastReroute   : Disabled
CSPF          : Disabled
Metric        : 0
Include Grps  :
None
Least Fill    : Disabled

LSP Tunnel ID : 1
Oper State     : Down
LSP Down Time  : 0d 00:00:07
Path Changes   : 0
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : 0
ClassType      : 0
Oper FR        : Disabled
ADSPEC         : Disabled
Exclude Grps   :
None

Auto BW        : Enabled
Auto BW Min    : 0 Mbps
AB Up Thresh   : 5 percent
AB Up BW       : 0 Mbps

AB OpState     : Down
Auto BW Max    : 100000 Mbps
AB Down Thresh : 5 percent
AB Down BW     : 0 Mbps

```

```

AB Curr BW : 0 Mbps
AB Adj Mul : 288+
AB Adj Time : 0 Mins
AB Adj Cnt : 0
AB Last Adj : n/a
ABMaxAvgRt : 0 Mbps
AB Ovfl Lmt : 0
ABOvflThres : 0 percent
AB Adj Cause: none
LdpOverRsvp : Enabled
IGP Shortcut: Enabled
Oper Metric : 65535

AB Samp Intv : 0
AB Samp Mul : 1+
AB Samp Time : 0 Mins
AB Samp Cnt : 0
AB Next Adj : 0 Mins
AB Lst AvgRt : 0 Mbps
AB Ovfl Cnt : 0
AB Ovfl BW : 0
AB Monitor BW : False
VprnAutoBind : Enabled

```

+ indicates inherited values

=====

*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls lsp count

=====

MPLS LSP Count

	Originate	Transit	Terminate
Static LSPs	0	136	0
Dynamic LSPs	140	421	1620
Detour LSPs	0	0	0
P2MP S2Ls	0	0	0

=====

*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls lsp path mbb

=====

MPLS LSP Paths

LSP to_110_20_1_1_cspf Path to_110_20_1_1

LastResignal: n/a	CSPF Metric : 1001
Last MBB :	
MBB Type : TimerBasedResignal	MBB State : Fail
Ended At : 03/04/2010 09:23:58	Old Metric : 0

LSP to_110_20_1_2_cspf Path to_110_20_1_2

LastResignal: 03/04/2010 09:23:58	CSPF Metric : 65535

LSP to_110_20_1_3_cspf Path to_110_20_1_3

LastResignal: n/a	CSPF Metric : 1001
Last MBB :	
MBB Type : TimerBasedResignal	MBB State : Fail
Ended At : 03/04/2010 09:23:58	Old Metric : 0

LSP to_110_20_1_4_cspf Path to_110_20_1_4

LastResignal: n/a	CSPF Metric : 0

```

LSP to_110_20_1_5_cspf Path to_110_20_1_5
-----
...
LastResignal: n/a                                CSPF Metric : 0
-----
LSP to_10_100_1_1_cspf_19 Path to_10_100_1_1
-----
LastResignal: n/a                                CSPF Metric : 0
-----
LSP to_10_100_1_1_cspf_20 Path to_10_100_1_1
-----
LastResignal: n/a                                CSPF Metric : 0
=====
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls lsp transit
=====
MPLS LSPs (Transit)
=====
Legend : @ - Active Detour
=====
From          To          In I/F    Out I/F    State LSP Name
-----
110.20.1.5    10.20.1.22    3/2/1     3/2/7     Up    to_10_20_1_22_cspf::to*
110.20.1.5    10.20.1.20    3/2/7     3/2/1     Up    to_10_20_1_20_cspf_3::*
110.20.1.5    10.20.1.20    3/2/7     3/2/1     Up    to_10_20_1_20_cspf_4::*
110.20.1.5    10.20.1.20    3/2/7     3/2/1     Up    to_10_20_1_20_cspf_2::*
110.20.1.5    10.20.1.20    3/2/7     3/2/1     Up    to_10_20_1_20_cspf_20:*
110.20.1.5    10.20.1.20    3/2/7     3/2/1     Up    to_10_20_1_20_cspf_18:*
110.20.1.5    10.20.1.20    3/2/7     3/2/1     Up    to_10_20_1_20_cspf_19:*
110.20.1.5    10.20.1.20    3/2/7     3/2/1     Up    to_10_20_1_20_cspf_17:*
110.20.1.5    10.20.1.20    3/2/7     3/2/1     Up    to_10_20_1_20_cspf_16:*
110.20.1.5    10.20.1.20    3/2/7     3/2/1     Up    to_10_20_1_20_cspf_15:*
110.20.1.5    10.20.1.20    3/2/7     3/2/1     Up    to_10_20_1_20_cspf_13:*
110.20.1.5    10.20.1.20    3/2/7     3/2/1     Up    to_10_20_1_20_cspf_14:*
110.20.1.5    10.20.1.20    3/2/7     3/2/1     Up    to_10_20_1_20_cspf_12:*
110.20.1.5    10.20.1.20    3/2/7     3/2/1     Up    to_10_20_1_20_cspf_10:*
...
110.20.1.3    10.20.1.22    aps-1     3/2/7     Up    to_10_20_1_22_cspf_6::*
110.20.1.3    10.20.1.22    aps-1     3/2/7     Up    to_10_20_1_22_cspf::to*
110.20.1.3    10.20.1.22    aps-1     3/2/7     Up    to_10_20_1_22_cspf_9::*
-----
LSPs : 520
=====
* indicates that the corresponding row element may have been truncated.
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls lsp terminate
=====
MPLS LSPs (Terminate)
=====
Legend : @ - Active Detour
=====
From          To          In I/F    Out I/F    State LSP Name
-----
110.20.1.5    110.20.1.4    3/2/1     n/a        Up    b4-1::b4-1
110.20.1.5    110.20.1.4    3/2/7     n/a        Up    gsr::gsr

```

```

10.20.1.22      110.20.1.4      3/2/7      n/a      Up      gsr2_t10
110.20.1.6      110.20.1.4      3/2/3:10   n/a      Up      1::2
110.20.1.6      110.20.1.4      3/2/3:3    n/a      Up      1::stby
110.20.1.6      110.20.1.4      3/2/3:10   n/a      Up      2::2
110.20.1.6      110.20.1.4      3/2/3:6    n/a      Up      2::stby
110.20.1.6      110.20.1.4      3/2/3:10   n/a      Up      3::2
110.20.1.6      110.20.1.4      3/2/3:6    n/a      Up      3::stby
...
110.20.1.3      110.20.1.4      aps-1      n/a      Up      to_110_20_1_4_cspf_20:*
110.20.1.3      110.20.1.4      aps-1      n/a      Up      to_110_20_1_4_cspf_4:.*
-----

```

LSPs : 1603

* indicates that the corresponding row element may have been truncated.

*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls lsp terminate detail

MPLS LSPs (Terminate) (Detail)

LSP b4-1::b4-1

```

-----
From          : 110.20.1.5          To          : 110.20.1.4
State         : Up
SetupPriority  : 7                  Hold Priority : 0
Class Type    : 0
In Interface  : 3/2/1              In Label    : 131071
Previous Hop  : 10.100.30.20
-----

```

LSP gsr::gsr

```

-----
From          : 110.20.1.5          To          : 110.20.1.4
State         : Up
SetupPriority  : 7                  Hold Priority : 0
Class Type    : 0
In Interface  : 3/2/7              In Label    : 128547
Previous Hop  : 160.60.60.2
-----

```

...

```

-----
From          : 110.20.1.3          To          : 110.20.1.4
State         : Up
SetupPriority  : 7                  Hold Priority : 0
Class Type    : 0
In Interface  : aps-1              In Label    : 130409
Previous Hop  : 104.104.0.3
-----

```

*A:SRU4>config>router>mpls#

*A:Dut-C>config>router>mpls>lsp# show router mpls lsp "N1.N2.3" detail

MPLS LSPs (Originating) (Detail)

Type : Originating

```

LSP Name       : N1.N2.3
LSP Type       : RegularLsp
From           : 1.0.0.1
Adm State      : Up
LSP Up Time    : 0d 00:01:07
Transitions    : 1
Retry Limit    : 0
Signaling      : RSVP
Hop Limit      : 255
Adaptive       : Enabled
FastReroute    : Disabled
CSPF           : Enabled
Metric         : 9
Load Balanc*   : 100
Include Grps   :
None
Least Fill     : Disabled

LSP Tunnel ID  : 132
To             : 1.0.0.2
Oper State     : Up
LSP Down Time  : 0d 00:00:00
Path Changes   : 2
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : 1500
ClassType      : 0
Oper FR        : Disabled
ADSPEC         : Disabled
Use TE metric   : Disabled

Exclude Grps   :
None

Revert Timer   : Disabled
Auto BW        : Disabled
LdpOverRsvp    : Enabled
IGP Shortcut   : Enabled
IGP LFA        : Disabled
BGPTransTun    : Enabled
Oper Metric    : 9
Prop Adm Grp   : Disabled

Next Revert In : N/A

VprnAutoBind   : Enabled
BGP Shortcut    : Enabled
IGP Rel Metric  : Disabled

Primary(a)      : path.N1.N2.3
Bandwidth       : 0 Mbps
Secondary       : path.N1.N2.4
Bandwidth       : 0 Mbps
Up Time         : 0d 00:01:04
Down Time       : 0d 00:01:00

```

```

=====
*A:SetupCLI# show router mpls lsp "lsp_1" path "500" detail

```

```

=====
MPLS LSP lsp_1 Path 500 (Detail)
=====

```

```

Legend :

```

```

@ - Detour Available          # - Detour In Use
b - Bandwidth Protected       n - Node Protected
s - Soft Preemption
S - Strict                    L - Loose
A - ABR

```

```

-----
LSP lsp_1 Path 500
-----

```

```

LSP Name       : lsp_1
Path LSP ID    : 38400
From           : 10.10.1.1
Admin State    : Up
Path Name      : 500
Path Admin     : Up
Out Interface   : n/a
Path Up Time    : 0d 00:00:00
Retry Limit    : 0
Retry Attempt   : 2
BFD Template    : None

To             : 10.10.1.2
Oper State     : Down
Path Type      : Primary
Path Oper      : Down
Out Label      : n/a
Path Down Time : 0d 00:00:43
Retry Timer    : 30 sec
Next Retry In  : 19 sec
BFD Ping Interval : 60

```

```

BFD Enable      : FALSE

Adspec          : Disabled      Oper Adspec      : N/A
CSPF            : Disabled      Oper CSPF       : N/A
Least Fill     : Disabled      Oper LeastFill  : N/A
FRR             : Disabled      Oper FRR        : N/A
Prop Adm Grp   : Disabled      Oper Prop AG    : N/A
Inter-area     : N/A

Neg MTU         : 0             Oper MTU        : N/A
Bandwidth      : No Reservation Oper Bw         : N/A
Hop Limit      : 255           Oper HopLimit    : N/A
Record Route   : Record        Oper Rec Route   : N/A
Record Label   : No Record     Oper Rec Label   : N/A
Setup Priority  : 7             Oper Setup Priority : N/A
Hold Priority   : 7             Oper Hold Priority : N/A
Class Type     : 7             Oper CT          : N/A
Backup CT      : 5
MainCT Retry   : 500
    Rem        :
MainCT Retry   : 500
    Limit      :
Include Grps   :               Oper InclGrps    :
None           :               N/A
Exclude Grps   :               Oper ExclGrps    :
None           :               N/A

Adaptive       : Enabled        Oper Metric      : N/A
Preference     : n/a
Path Trans     : 0              CSPF Queries     : 0
Failure Code   : noResourcesAvailable
Failure Node   : 9.1.255.255
Explicit Hops   :
    No Hops Specified
Actual Hops     :
    No Hops Specified
Resignal Eligible: False
Last Resignal  : n/a           CSPF Metric       : N/A
=====

```

*A:Dut-B# show router mpls lsp path detail tof1

=====

MPLS LSP tof1 Path (Detail)

=====

Legend :

```

@ - Detour Available      # - Detour In Use
b - Bandwidth Protected   n - Node Protected
s - Soft Preemption

```

LSP tof1 Path 1

```

LSP Name      : tof1                Path LSP ID : 4706
From          : 10.20.1.2           To          : 10.20.1.4
Adm State     : Up                  Oper State   : Up
Path Name     : 1                   Path Type    : Primary
Path Admin    : Up                  Path Oper    : Up
OutInterface  : 1/1/2:5             Out Label    : 124809
Path Up Time  : 0d 00:01:16         Path Dn Time : 0d 00:00:00

```

```

Retry Limit : 0
RetryAttempt: 0
SetupPriorit*: 4
Bandwidth : 1 Mbps
Hop Limit : 255
Record Route: Record
Oper MTU : 1492
Adaptive : Enabled
Include Grps:
None
Path Trans : 2
Failure Code: noError
ExplicitHops:
  10.20.1.4
Actual Hops :
  10.10.105.2(10.20.1.2) @ s
  -> 10.10.105.4(10.20.1.4)
ComputedHops:
  10.10.105.2 -> 10.10.105.4
ResigEligib*: False
LastResignal: n/a
In Prog MBB :
  MBB Type : SoftPreemption
  Started At : 12/08/2008 22:21:11
  FailureCode: noError
Retry Timer : 20 sec
NextRetryIn : 0 sec
Hold Priorit*: 4
Oper Bw : 1 Mbps
Class Type : 5
Record Label: Record
Neg MTU : 1492
Oper Metric : 1000
Exclude Grps:
None
CSPF Queries: 1
Failure Node: n/a
Record Label : N/A
Record Label : 124809
CSPF Metric : 1000
NextRetryIn : 19 sec
RetryAttempt: 0
Failure Node: n/a

```

=====

*A:Dut-B#

*A:PE-4>config>router>mpls# show router mpls lsp "LSP-PE-2-PE-4" detail

=====

MPLS LSPs (Originating) (Detail)

Type : Originating

```

-----
LSP Name      : LSP-PE-2-PE-4
LSP Type      : RegularLsp
LSP Index     : 4
From          : 192.0.2.104
Adm State     : Up
LSP Up Time   : 17d 02:16:53
Transitions   : 3
Retry Limit   : 0
Signaling     : RSVP
Hop Limit     : 255
Adaptive      : Enabled
FastReroute   : Disabled
CSPF          : Disabled
Metric        : N/A
Load Bal Wt   : N/A
Include Grps:
None
Least Fill    : Disabled
BFD Template  : None
BFD Enable    : False
LSP Tunnel ID : 4
TTM Tunnel Id : 4
To            : 192.0.2.102
Oper State    : Up
LSP Down Time : 0d 00:00:00
Path Changes  : 3
Retry Timer   : 30 sec
Resv. Style   : SE
Negotiated MTU : 8690
ClassType     : 0
Oper FR       : Disabled
ADSPEC        : Disabled
ClassForwarding: Disabled
Exclude Grps  :
None
BFD Ping Intvl : 60
Next Revert In : N/A
AB OpState    : Down

```



```

Auto BW Min : 0 Mbps
AB Up Thresh: 5 percent
AB Up BW : 0 Mbps
AB Curr BW : 0 Mbps
AB Adj Mul : 288+
AB Adj Time : 0 Mins
AB Adj Cnt : 0
AB Last Adj : n/a
ABMaxAvgRt : 0 Mbps
AB Ovfl Lmt : 0
ABOvflThres : 0 percent
AB UndflLmt : 0
ABUndflThrs : 0 percent
ABMaxUndflBW: 0 Mbps
AB Adj Cause: none
Be Weight : 100 percent
L1 Weight : 100 percent
Nc Weight : 100 percent
H1 Weight : 100 percent
LdpOverRsvp : Enabled
VprnAutoBind: Enabled
IGP Shortcut: Enabled
IGP LFA : Disabled
BGPTransTun : Enabled
Oper Metric : 1
Prop Adm Grp: Disabled

Auto BW Max : 100000 Mbps
AB Down Thresh : 5 percent
AB Down BW : 0 Mbps
AB Samp Intv : 0 Mins
AB Samp Mul : 1+
AB Samp Time : 0 Mins
AB Samp Cnt : 0
AB Next Adj : 0 Mins
AB Lst AvgRt : 0 Mbps
AB Ovfl Cnt : 0
AB Ovfl BW : 0 Mbps
AB Undrfl Cnt : 0
AB Undrfl BW : 0 Mbps

AB Monitor BW : False
Af Weight : 100 percent
L2 Weight : 100 percent
Ef Weight : 100 percent
H2 Weight : 100 percent

BGP Shortcut : Enabled
IGP Rel Metric : Disabled

Primary(a) : completely-loose-path
Bandwidth : 0 Mbps

Up Time : 17d 02:16:53
=====

```

lsp-egress-stats

Syntax **lsp-egress-stats**
lsp-egress-stats *lsp-name*

Context show>router>mpls

Description This command displays MPLS LSP egress statistics information.

Output

Sample Output

```

*A:Dut-C>config>router>mpls>lsp$ show router mpls lsp-egress-stats lsp "1"

=====
MPLS LSP Egress Statistics
=====
-----
LSP Name      : 1
-----
Collect Stats : Enabled
Adm State     : Up
FC BE

Accting Plcy. : Default
PSB Match     : True

```

```

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
=====

```

```
*A:Dut-C# show router mpls lsp-egress-stats lsp "ipmsi-1-73728"
```

```
=====
MPLS LSP Egress Statistics
=====
```

```
-----
LSP Name      : ipmsi-1-73728
-----
```

```

Collect Stats : Enabled                      Accting Plcy. : Default
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-ingress-stats

Syntax	lsp-ingress-stats lsp-ingress-stats <i>ip-address</i> lsp <i>lsp-name</i>
Context	show>router>mpls
Description	This command displays MPLS LSP ingress statistics information.
Output	

Sample Output

```
*A:Dut-A# show router mpls lsp-ingress-stats lsp "1" sender 10.20.1.3

=====
MPLS LSP Ingress Statistics
=====
-----
LSP Name      : 1
Sender        : 10.20.1.3
-----
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
=====

*A:Dut-A# show router mpls lsp-ingress-stats lsp "ipmsi-1-73728" sender 10.20.1.3
```

```
=====
MPLS LSP Ingress Statistics
=====
-----
LSP Name       : ipmsi-1-73728
Sender         : 10.20.1.3
-----
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
=====
```

```
*A:Dut-A>config>router>mpls>ingr-stats# show router mpls lsp-ingress-stats
type p2mp active template-match
```

```
=====
MPLS LSP Ingress Statistics
=====
-----
LSP Name       : ipmsi-1-73728
Sender         : 10.20.1.3
-----
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 Statistics : 1

```

lsp-template

Syntax **lsp-template** [*/lsp-template-name*] [**detail**]

Context show>router>mpls

Description This command displays MPLS LSP template information.

Output

Sample Output

```

*A:Dut-C# show router mpls lsp-template detail
=====
MPLS LSP Templates (Detail)
=====
-----
LSP Template : ipmsi
-----
Type           : P2MP                      Admin State    : Up
Default Path    : path_ipmsi              Adaptive      : Enabled
Bandwidth       : 0 Mbps                   Hop Limit      : 255
CSPF            : Enabled                  Use TE metric   : Disabled
Include Groups  :                          Exclude Groups  :
None                                                    None
FastReroute     : Enabled
FR Method       : Facility                 FR Hop Limit   : 16
Record Route    : Record                  Record Label    : Record
Retry Limit     : 0                      Retry Timer     : 30 sec
LSP Count       : 3                      Ref Count      : 3
=====

```

oam-template

Syntax **oam-template**

Context show>router>mpls>mpls-tp

Description This command displays MPLS-TP OAM template information.

Output**Sample Output**

```
*A:mlstp-dutA# show router mpls mpls-tp oam-template
```

```
=====
MPLS-TP OAM Templates
=====
```

```
Template Name : privatebed-oam-template Router ID      : 1
BFD Template  : privatebed-bfd-template Hold-Down Time: 0 centiseconds
                                           Hold-Up Time  : 20 deciseconds
=====
```

protection-template

Syntax **protection-template**

Context show>router>mpls>mpls-tp

Description This command displays MPLS-TP protection template information.

Output**Sample Output**

```
*A:mlstp-dutA# show router mpls mpls-tp protection-template
```

```
=====
MPLS-TP Protection Templates
=====
```

```
Template Name : privatebed-protection-template Router ID      : 1
Protection Mode: one2one Direction              : bidirectional
Revertive      : revertive Wait-to-Restore: 300sec
Rapid-PSC-Timer: 10ms Slow-PSC-Timer : 5sec
=====
```

status

Syntax **status**

Context show>router>mpls>mpls-tp

Description This command displays MPLS-TP system configuration information.

Output

Sample Output

```
*A:mlstp-dutA# show router mpls mpls-tp status

=====
MPLS-TP Status
=====
Admin Status   : Up
Global ID      : 42
Node ID        : 0.0.3.233
Tunnel Id Min  : 1
Tunnel Id Max  : 4096
=====
```

transit-path

Syntax	transit-path [<i>path-name</i>] [<i>detail</i>]
Context	show>router>mpls>mpls-tp
Description	This command displays MPLS-TP tunnel information.
Parameters	<i>path-name</i> — Specifies the path name, up to 32 characters max.
Output	

Sample Output

```
A:mlstp-dutC# show router mpls mpls-tp transit-path
<path-name>
"tp-32"  "tp-33"  "tp-34"  "tp-35"  "tp-36"  "tp-37"  "tp-38"  "tp-39"
"tp-40"  "tp-41"
detail

A:mlstp-dutC# show router mpls mpls-tp transit-path "tp-32"

=====
MPLS-TP Transit tp-32 Path Information
=====
Path Name      : tp-32
Admin State    : Up
Oper State     : Up

-----
Path           NextHop           InLabel  OutLabel  Out I/F
-----
FP             2080             2081     CtoB_1
RP             2081             2080     CtoA_1
=====
A:mlstp-dutC# show router mpls mpls-tp transit-path "tp-32" detail

=====
MPLS-TP Transit tp-32 Path Information (Detail)
=====
Path Name      : tp-32
Admin State    : Up
Oper State     : Up
-----
Path ID configuration
```

```
Src Global ID : 42                               Dst Global ID : 42
Src Node ID   : 0.0.3.234                         Dst Node ID   : 0.0.3.233
LSP Number    : 2                                 Dst Tunnel Num: 32

Forward Path configuration
In Label      : 2080                               Out Label     : 2081
Out Interface : CtoB_1                             Next Hop Addr : n/a

Reverse Path configuration
In Label      : 2081                               Out Label     : 2080
Out Interface : CtoA_1                             Next Hop Addr : n/a
=====
A:mplstp-dutC#
```

p2mp-info

- Syntax**
- p2mp-info** [**type** {**originate** | **transit** | **terminate**}] [**s2l-endpoint** *ip-address*]
- Context**
- show>router>mpls
- Description**
- This command displays P2MP cross-connect information.
- Parameters**
- type** — Specifies the P2MP type.

Values

originate — Specifies to display the static LSPs that originate at this virtual router.

transit — Specifies to display the static LSPs that transit through this virtual router.

terminate — Specifies to display the static LSPs that terminate at this virtual router.

Output

Sample Output

```
*A:Dut-C# show router mpls p2mp-info
=====
MPLS P2MP Cross Connect Information
=====

-----
S2L ipmsi-4000-73729::path_ipmsi
-----
Source IP Address   : 10.20.1.1                Tunnel ID   : 61441
P2MP ID            : 4000                      Lsp ID     : 29696
S2L Name           : ipmsi-4000-73729::pa*      To        : 10.20.1.3
In Interface       : 1/1/1                     In Label   : 262129
Num. of S2ls      : 1
-----

S2L ipmsi-65535-73730::path_ipmsi
-----
Source IP Address   : 10.20.1.1                Tunnel ID   : 61442
P2MP ID            : 65535                      Lsp ID     : 30208
S2L Name           : ipmsi-65535-73730::p*      To        : 10.20.1.3
```



```

In Interface      : 1/1/1                In Label      : 262128
Num. of S2ls     : 1
-----
S2L ipmsi-1001-73728::path_ipmsi
-----
Source IP Address : 10.20.1.1            Tunnel ID    : 61440
P2MP ID          : 1001                  Lsp ID      : 35840
S2L Name         : ipmsi-1001-73728::pa* To      : 10.20.1.3
In Interface     : 1/1/1                In Label    : 262127
Num. of S2ls     : 1
-----
S2L ipmsi-1001-73732::path_ipmsi
-----
Source IP Address : 10.20.1.2            Tunnel ID    : 64944
P2MP ID          : 1001                  Lsp ID      : 34816
S2L Name         : ipmsi-1001-73732::pa* To      : 10.20.1.3
In Interface     : 1/1/2                In Label    : 262114
Num. of S2ls     : 1
-----
S2L ipmsi-4000-73729::path_ipmsi
-----
Source IP Address : 10.20.1.3            Tunnel ID    : 61441
P2MP ID          : 4000                  Lsp ID      : 16384
S2L Name         : ipmsi-4000-73729::pa* To      : 10.20.1.1
Out Interface    : 1/1/1                Out Label   : 262131
Num. of S2ls     : 1
-----
S2L ipmsi-4000-73729::path_ipmsi
-----
Source IP Address : 10.20.1.3            Tunnel ID    : 61441
P2MP ID          : 4000                  Lsp ID      : 16384
S2L Name         : ipmsi-4000-73729::pa* To      : 10.20.1.4
Out Interface    : 2/1/1                Out Label   : 262121
Num. of S2ls     : 1
-----
S2L ipmsi-1001-73728::path_ipmsi
-----
Source IP Address : 10.20.1.3            Tunnel ID    : 61440
P2MP ID          : 1001                  Lsp ID      : 22016
S2L Name         : ipmsi-1001-73728::pa* To      : 10.20.1.1
Out Interface    : 1/1/1                Out Label   : 262129
Num. of S2ls     : 1
-----
S2L ipmsi-1001-73728::path_ipmsi
-----
Source IP Address : 10.20.1.3            Tunnel ID    : 61440
P2MP ID          : 1001                  Lsp ID      : 22016
S2L Name         : ipmsi-1001-73728::pa* To      : 10.20.1.2
Out Interface    : 1/1/2                Out Label   : 262115
Num. of S2ls     : 1
-----
S2L ipmsi-1001-73728::path_ipmsi
-----
Source IP Address : 10.20.1.3            Tunnel ID    : 61440
P2MP ID          : 1001                  Lsp ID      : 22016
S2L Name         : ipmsi-1001-73728::pa* To      : 10.20.1.4
Out Interface    : 2/1/1                Out Label   : 262108
Num. of S2ls     : 2
-----

```

```

S2L ipmsi-1001-73728::path_ipmsi
-----
Source IP Address      : 10.20.1.3          Tunnel ID      : 61440
P2MP ID                : 1001              Lsp ID        : 22016
S2L Name               : ipmsi-1001-73728::pa* To      : 10.20.1.5
Out Interface          : 2/1/1             Out Label     : 262108
Num. of S2ls           : 2
-----

S2L ipmsi-65535-73730::path_ipmsi
-----
Source IP Address      : 10.20.1.3          Tunnel ID      : 61442
P2MP ID                : 65535             Lsp ID        : 46592
S2L Name               : ipmsi-65535-73730::p* To      : 10.20.1.1
Out Interface          : 1/1/1             Out Label     : 262130
Num. of S2ls           : 1
-----

S2L ipmsi-65535-73730::path_ipmsi
-----
Source IP Address      : 10.20.1.3          Tunnel ID      : 61442
P2MP ID                : 65535             Lsp ID        : 46592
S2L Name               : ipmsi-65535-73730::p* To      : 10.20.1.4
Out Interface          : 2/1/1             Out Label     : 262109
Num. of S2ls           : 1
-----

P2MP Cross-connect instances : 12

```

p2mp-lsp

Syntax **p2mp-lsp** [*/lsp-name*] [**detail**]
p2mp-lsp [*/lsp-name*] **p2mp-instance** [*p2mp-instance-name*] [**mbb**]
p2mp-lsp [*/lsp-name*] **p2mp-instance** [*p2mp-instance-name*] **s2l** [*s2l-name*] [**to** *s2l-to-address*]] [**status** {**up** | **down**}] [**detail**]
p2mp-lsp [*/lsp-name*] **p2mp-instance** [*p2mp-instance-name*] **s2l** [*s2l-name*] [**to** *s2l-to-address*]] **mbb**

Context show>router>mpls

Description This command displays MPLS P2MP LSP information.

Parameters *lsp-name* — Specifies the name of the LSP used in the path.
p2mp-instance[*p2mp-instance-name* — Specifies the administrative name for the P2MP instance which must be unique within a virtual router instance.
mbb — Specifies to display make-before-break (MBB) information.
s2l — Specifies the source-to-leaf (S2L) name.
to *s2l-to-address* — Indicates the IP address of the destination address of the S2L sub-LSP.

status — Displays the status of the p2mp LSP.

Values	up — Displays the total time that this S2I has been operational. down — Displays the total time that this S2I has not been operational.

Output

Sample Output

```
*A:Dut-C# show router mpls p2mp-lsp
- p2mp-lsp [<lsp-name>] [detail]
- p2mp-lsp [<lsp-name>] p2mp-instance [<p2mp-instance-name>] [mbb]
- p2mp-lsp [<lsp-name>] p2mp-instance [<p2mp-instance-name>] s2l [<s2l-name>]
  [to <s2l-to-address>]] [status {up | down}] [detail]
- p2mp-lsp [<lsp-name>] p2mp-instance [<p2mp-instance-name>] s2l [<s2l-name>]
  [to <s2l-to-address>]] <mbb>
- p2mp-lsp using-template [lsp-template <template-name>] [detail]

<lsp-name>          : [64 chars max] - accepts * as wildcard char
<p2mp-instance>     : keyword
<p2mp-instance-name> : [max 32 chars]
<s2l>               : keyword
<s2l-name>          : [max 32 chars]
<up | down>         : keywords
<detail>            : keyword
<mbb>               : keyword
<s2l-to-address>    : [a.b.c.d]
<using-template>    : keyword
<lsp-template>     : [32 chars max]
```

```
*A:Dut-C# show router mpls p2mp-lsp

=====
MPLS P2MP LSPs (Originating)
=====
LSP Name                                                    Adm   Opr
-----
ipmsi-1001-73728                                           Up    Up
ipmsi-4000-73729                                           Up    Up
ipmsi-65535-73730                                           Up    Up
-----

LSPs : 3

=====
*A:Dut-C# show router mpls p2mp-lsp detail

=====
MPLS P2MP LSPs (Originating) (Detail)
=====
Type : Originating
-----
LSP Name      : ipmsi-1001-73728
```

```

LSP Type       : P2mpAutoLsp
From           : 10.20.1.3
Adm State      : Up
LSP Up Time    : 6d 21:08:37
Transitions    : 1
Retry Limit    : 0
Signaling      : RSVP
Hop Limit      : 255
Adaptive       : Enabled
FastReroute    : Enabled
FR Method      : Facility
FR Bandwidth   : 0 Mbps
FR Object      : Enabled
CSPF           : Enabled
Metric         : Disabled
Include Grps   :
None
Least Fill     : Disabled

LSP Tunnel ID  : 61440

Oper State     : Up
LSP Down Time  : 0d 00:00:00
Path Changes   : 1
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : n/a
ClassType      : 0
Oper FR        : Enabled
FR Hop Limit   : 16
FR Node Protect : Disabled

ADSPEC         : Disabled
Use TE metric  : Disabled
Exclude Grps   :
None

Auto BW        : Disabled
LdpOverRsvp    : Disabled
IGP Shortcut    : Disabled
BGPTransTun    : Disabled
Oper Metric    : Disabled
Prop Adm Grp   : Disabled
P2MPInstance   : 1001
S2L Cfg Cou*   : 4
S2L-Name       : path_ipmsi
S2L-Name       : path_ipmsi
S2L-Name       : path_ipmsi
S2L-Name       : path_ipmsi

VprnAutoBind   : Disabled
BGP Shortcut    : Disabled

CSPFFirstLoose : Disabled
P2MP-Inst-type : Primary
S2L Oper Count* : 4
To              : 10.20.1.1
To              : 10.20.1.2
To              : 10.20.1.4
To              : 10.20.1.5

-----
Type : Originating
-----

LSP Name       : ipmsi-4000-73729
LSP Type       : P2mpAutoLsp
From           : 10.20.1.3
Adm State      : Up
LSP Up Time    : 6d 21:08:38
Transitions    : 1
Retry Limit    : 0
Signaling      : RSVP
Hop Limit      : 255
Adaptive       : Enabled
FastReroute    : Enabled
FR Method      : Facility
FR Bandwidth   : 0 Mbps
FR Object      : Enabled
CSPF           : Enabled
Metric         : Disabled
Include Grps   :
None
Least Fill     : Disabled

LSP Tunnel ID  : 61441

Oper State     : Up
LSP Down Time  : 0d 00:00:00
Path Changes   : 1
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : n/a
ClassType      : 0
Oper FR        : Enabled
FR Hop Limit   : 16
FR Node Protect : Disabled

ADSPEC         : Disabled
Use TE metric  : Disabled
Exclude Grps   :
None

Auto BW        : Disabled
LdpOverRsvp    : Disabled
IGP Shortcut    : Disabled
BGPTransTun    : Disabled

VprnAutoBind   : Disabled
BGP Shortcut    : Disabled

```

```

Oper Metric : Disabled
Prop Adm Grp: Disabled
CSPFFirstLoose : Disabled

P2MPInstance: 4000
S2L Cfg Cou*: 2
S2l-Name : path_ipmsi
S2l-Name : path_ipmsi
P2MP-Inst-type : Primary
S2L Oper Count*: 2
To : 10.20.1.1
To : 10.20.1.4
-----
Type : Originating
-----
LSP Name : ipmsi-65535-73730
LSP Type : P2mpAutoLsp
From : 10.20.1.3
LSP Tunnel ID : 61442
Adm State : Up
Oper State : Up
LSP Up Time : 6d 21:08:39
LSP Down Time : 0d 00:00:00
Transitions : 1
Path Changes : 1
Retry Limit : 0
Retry Timer : 30 sec
Signaling : RSVP
Resv. Style : SE
Hop Limit : 255
Negotiated MTU : n/a
Adaptive : Enabled
ClassType : 0
FastReroute : Enabled
Oper FR : Enabled
FR Method : Facility
FR Hop Limit : 16
FR Bandwidth: 0 Mbps
FR Node Protect: Disabled
FR Object : Enabled
CSPF : Enabled
ADSPEC : Disabled
Metric : Disabled
Use TE metric : Disabled
Include Grps:
Exclude Grps :
None
None
Least Fill : Disabled

Auto BW : Disabled
LdpOverRsvp : Disabled
VprnAutoBind : Disabled
IGP Shortcut: Disabled
BGP Shortcut : Disabled
BGPTransTun : Disabled
Oper Metric : Disabled
Prop Adm Grp: Disabled
CSPFFirstLoose : Disabled

P2MPInstance: 65535
S2L Cfg Cou*: 2
S2l-Name : path_ipmsi
S2l-Name : path_ipmsi
P2MP-Inst-type : Primary
S2L Oper Count*: 2
To : 10.20.1.1
To : 10.20.1.4
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-C#

```

*A:sim1>config>router>mpls>lsp\$ show router mpls p2mp-lsp p2mp-instance s2l detail

```

=====
MPLS LSP S2L (Detail)
=====
Legend :
  @ - Detour Available          # - Detour In Use
  b - Bandwidth Protected      n - Node Protected
  S - Strict                   L - Loose
  s - Soft Preemption
=====
-----
LSP 1 S2L 1

```

```

-----
LSP Name      : 1
P2MP ID       : 0
Adm State     : Up
S2l State:    : Inactive
S2L Name      : 1
S2l Admin     : Up
OutInterface:  n/a
S2L Up Time   : 0d 00:00:00
RetryAttempt:  0
S2L Trans     : 8
Failure Code:  noError
ExplicitHops:
    10.20.1.2(S)
Actual Hops   :
    No Hops Specified
ComputedHops:
    No Hops Specified
LastResignal: n/a
=====

```

```
show router mpls p2mp-lsp p2mp-instance s2l detail
```

```

-----
LSP 2 S2L 2
-----
LSP Name      : 2
P2MP ID       : 0
Adm State     : Up
S2l State:    : Active
S2L Name      : 2
S2l Admin     : Up
OutInterface:  1/1/1
S2L Up Time   : 0d 00:04:43
RetryAttempt:  0
S2L Trans     : 5
Failure Code:  tunnelLocallyRepaired
ExplicitHops:
    10.20.1.2(S)
Actual Hops   :
    10.10.1.1(10.20.1.1)
    -> 10.10.1.2(10.20.1.2) @ #
    -> 10.10.6.3(10.20.1.3)
ComputedHops:
    10.10.1.1(S)      -> 10.10.1.2(S)      -> 10.10.5.3(S)
LastResignal: n/a
In Prog MBB   :
    MBB Type      : GlobalRevert
    Timeout In    : 23 sec
    Started At    : 06/29/2011 11:06:09
    FailureCode:  noError
S2l LSP ID    : 52230
S2l Grp Id    : 2
Oper State    : Up
To            : 10.20.1.3
S2l Oper      : Up
Out Label     : 131071
S2L Dn Time   : 0d 00:00:00
NextRetryIn   : 0 sec
CSPF Queries: 21
Failure Node:  10.20.1.2
Record Label   : N/A
Record Label   : 131071
Record Label   : 131068
NextRetryIn    : n/a
RetryAttempt:  7
Failure Node:  n/a
=====

```

```
*A:Dut-C>config>router>mpls>lsp$ /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
=====
-----
LSP 1 Path 1
-----
LSP Name       : 1                      Path LSP ID : 56320
From           : 10.20.1.3              To           : 10.10.1.1
Adm State      : Up                     Oper State   : Up
Path Name      : 1                      Path Type    : Primary
Path Admin     : Up                     Path Oper    : Up
OutInterface   : 1/1/1                  Out Label    : 131071
Path Up Time   : 0d 00:03:09            Path Dn Time : 0d 00:00:00
Retry Limit    : 0                      Retry Timer   : 30 sec
RetryAttempt   : 0                      NextRetryIn  : 0 sec
SetupPrioriti* : 7                      Hold Prioriti* : 0
Preference     : n/a
Bandwidth      : No Reservation          Oper Bw       : 0 Mbps
Hop Limit      : 255                    Class Type    : 0
Backup CT      : None
MainCT Retry   : n/a                    MainCT Retry  : 0
Rem            :                          Limit         :
Oper CT        : 0
Record Route   : Record                  Record Label  : Record
Oper MTU       : 1496                    Neg MTU       : 1496
Adaptive       : Enabled                  Oper Metric   : 1000
Include Grps   :                          Exclude Grps  :
None                                                    None
Path Trans     : 1                      CSPF Queries  : 3
Failure Code   : badNode                  Failure Node   : 10.20.1.3

Oper Values :
Setup Prior* : 7                      Hold Prioriti* : 0
Record Route : Record                  Record Label   : Record
Hop Limit    : 255
Adspec       : Disabled
CSPF         : Enabled                  CSPFToFirst*   : Disabled
Least Fill   : Disabled                  FR Node Pro*   : Disabled
Prop Adm Grp : Disabled
Include Grps :                          Exclude Grps   :
None                                                    None

ExplicitHops:
  No Hops Specified
Actual Hops :
  10.10.2.3(10.20.1.3) @ #              Record Label   : N/A
  -> 10.10.1.1(10.20.1.1)                Record Label   : 131071
ComputedHops:
  10.10.2.3(S) -> 10.10.2.1(S)
ResigEligib* : False
LastResignal  : n/a                      CSPF Metric    : 1000
In Prog MBB   :
  MBB Type    : GlobalRevert              NextRetryIn    : 0 sec
  Timeout In  : 22 sec
  Started At  : 08/26/2011 23:59:29        RetryAttempt    : 2

```

```

FailureCode: noError                      Failure Node: n/a
Signaled BW: 0 Mbps
=====
* indicates that the corresponding row element may have been truncated.

show router mpls p2mp-lsp p2mp-instance s2l detail
-----
LSP 2 S2L 2
-----
LSP Name       : 2                      S2L LSP ID    : 52230
P2MP ID        : 0                      S2L Grp Id   : 4
Adm State      : Up                     Oper State    : Down
S2L State:     : Inactive                :
S2L Name       : 2                      To            : 10.20.1.3
S2L Admin      : Up                     S2L Oper      : In Progress
OutInterface:  n/a                     Out Label     : n/a
S2L Up Time    : 0d 00:00:00            S2L Dn Time   : 0d 00:00:20
RetryAttempt:  1                      NextRetryIn   : n/a
Timeout In     : 21 sec
S2L Trans      : 6                      CSPF Queries  : 27
Failure Code:  noError                 Failure Node:  n/a
ExplicitHops:
    10.20.1.2(S)
Actual Hops    :
    No Hops Specified
LastResignal:  n/a

```

```

*A:Dut-C# show router mpls p2mp-lsp
=====
MPLS P2MP LSPs (Originating)
=====
LSP Name                               Adm   Opr
-----
ipmsi-1001-73728                       Up    Up
ipmsi-4000-73729                       Up    Up
ipmsi-65535-73730                      Up    Up
-----
LSPs : 3
=====

```

```

*A:Dut-C# show router mpls p2mp-lsp detail
=====
MPLS P2MP LSPs (Originating) (Detail)
=====
Type : Originating
-----
LSP Name       : ipmsi-1001-73728
LSP Type       : P2mpAutoLsp          LSP Tunnel ID : 61440
From           : 10.20.1.3
Adm State      : Up                     Oper State     : Up
LSP Up Time    : 6d 21:08:37            LSP Down Time : 0d 00:00:00
Transitions    : 1                      Path Changes   : 1
Retry Limit    : 0                      Retry Timer    : 30 sec
Signaling      : RSVP                   Resv. Style    : SE
Hop Limit      : 255                    Negotiated MTU : n/a

```



```

Adaptive      : Enabled
FastReroute   : Enabled
FR Method     : Facility
FR Bandwidth  : 0 Mbps
FR Object     : Enabled
CSPF          : Enabled
Metric        : Disabled
Include Grps  :
None
Least Fill    : Disabled
Auto BW       : Disabled
LdpOverRsvp   : Disabled
IGP Shortcut  : Disabled
BGPTransTun   : Disabled
Oper Metric   : Disabled
Prop Adm Grp  : Disabled
P2MPInstance  : 1001
S2L Cfg Cou*  : 4
S2l-Name      : path_ipmsi
S2l-Name      : path_ipmsi
S2l-Name      : path_ipmsi
S2l-Name      : path_ipmsi
-----
Type : Originating
-----
LSP Name      : ipmsi-4000-73729
LSP Type      : P2mpAutoLsp
From          : 10.20.1.3
Adm State     : Up
LSP Up Time   : 6d 21:08:38
Transitions   : 1
Retry Limit   : 0
Signaling     : RSVP
Hop Limit     : 255
Adaptive      : Enabled
FastReroute   : Enabled
FR Method     : Facility
FR Bandwidth  : 0 Mbps
FR Object     : Enabled
CSPF          : Enabled
Metric        : Disabled
Include Grps  :
None
Least Fill    : Disabled
Auto BW       : Disabled
LdpOverRsvp   : Disabled
IGP Shortcut  : Disabled
BGPTransTun   : Disabled
Oper Metric   : Disabled
Prop Adm Grp  : Disabled
P2MPInstance  : 4000
S2L Cfg Cou*  : 2
S2l-Name      : path_ipmsi
S2l-Name      : path_ipmsi
-----
Type : Originating
-----
LSP Name      : ipmsi-65535-73730
LSP Type      : P2mpAutoLsp
-----
ClassType     : 0
Oper FR       : Enabled
FR Hop Limit  : 16
FR Node Protect : Disabled
ADSPEC        : Disabled
Use TE metric : Disabled
Exclude Grps  :
None
VprnAutoBind  : Disabled
BGP Shortcut  : Disabled
CSPFFirstLoose : Disabled
P2MP-Inst-type : Primary
S2L Oper Count* : 4
To            : 10.20.1.1
To            : 10.20.1.2
To            : 10.20.1.4
To            : 10.20.1.5
-----
LSP Tunnel ID : 61441
Oper State     : Up
LSP Down Time : 0d 00:00:00
Path Changes   : 1
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : n/a
ClassType     : 0
Oper FR       : Enabled
FR Hop Limit  : 16
FR Node Protect : Disabled
ADSPEC        : Disabled
Use TE metric : Disabled
Exclude Grps  :
None
VprnAutoBind  : Disabled
BGP Shortcut  : Disabled
CSPFFirstLoose : Disabled
P2MP-Inst-type : Primary
S2L Oper Count* : 2
To            : 10.20.1.1
To            : 10.20.1.4
-----

```

```

From       : 10.20.1.3
Adm State  : Up
LSP Up Time : 6d 21:08:39
Transitions : 1
Retry Limit : 0
Signaling   : RSVP
Hop Limit   : 255
Adaptive    : Enabled
FastReroute : Enabled
FR Method    : Facility
FR Bandwidth: 0 Mbps
FR Object    : Enabled
CSPF        : Enabled
Metric      : Disabled
Include Grps:
None
Least Fill  : Disabled
Auto BW     : Disabled
LdpOverRsvp : Disabled
IGP Shortcut: Disabled
BGPTransTun : Disabled
Oper Metric : Disabled
Prop Adm Grp: Disabled
P2MPInstance: 65535
S2L Cfg Cou*: 2
S2l-Name    : path_ipmsi
S2l-Name    : path_ipmsi

Oper State  : Up
LSP Down Time : 0d 00:00:00
Path Changes  : 1
Retry Timer   : 30 sec
Resv. Style   : SE
Negotiated MTU : n/a
ClassType     : 0
Oper FR       : Enabled
FR Hop Limit  : 16
FR Node Protect: Disabled

ADSPEC       : Disabled
Use TE metric : Disabled
Exclude Grps :
None

VprnAutoBind : Disabled
BGP Shortcut  : Disabled

CSPFFirstLoose : Disabled
P2MP-Inst-type : Primary
S2L Oper Count*: 2
To              : 10.20.1.1
To              : 10.20.1.4

```

=====

* indicates that the corresponding row element may have been truncated.

*A:Dut-C#

*A:sim1>config>router>mpls>lsp\$ show router mpls p2mp-lsp p2mp-instance s2l detail

=====

MPLS LSP S2L (Detail)

Legend :

@ - Detour Available	# - Detour In Use
b - Bandwidth Protected	n - Node Protected
S - Strict	L - Loose
s - Soft Preemption	

LSP 1 S2L 1

```

LSP Name      : 1
P2MP ID       : 0
Adm State     : Up
S2l State:    : Inactive
S2L Name      : 1
S2l Admin     : Up
OutInterface:  : n/a
S2L Up Time   : 0d 00:00:00
RetryAttempt:  : 0
S2L Trans     : 8
Failure Code:  : noError
ExplicitHops:  :
    10.20.1.2(S)
Actual Hops   :

S2l LSP ID    : 26624
S2l Grp Id    : 0
Oper State    : Down
To            : 10.20.1.3
S2l Oper      : Down
Out Label     : n/a
S2L Dn Time   : 0d 00:00:01
NextRetryIn   : 9 sec (Fast)
CSPF Queries: : 4
Failure Node:  : n/a

```

```

    No Hops Specified
ComputedHops:
    No Hops Specified
LastResignal: n/a
=====

A:ALU-25# show router mpls p2mp-lsp lsp_1
=====
MPLS LSPs (Originating)
=====
LSP Name                                To/P2MP ID                Fastfail    Adm   Opr
                                Config
-----
lsp_1                                18                        Yes         Up    Up
-----
LSPs : 1
=====
A:ALU-25#

A:ALU-25# show router mpls p2mp-lsp Test_p2mp detail
=====
MPLS P2MP LSPs (Originating) (Detail)
=====
Type : Originating
-----
LSP Name      : lsp_1                LSP Tunnel ID : 1
From          : 10.10.1.1            P2MP ID       : 18

Adm State     : Up                  Oper State     : Down
LSP Up Time   : 0d 00:00:00         LSP Down Time : 0d 20:39:48
Transitions   : 0                  Path Changes   : 0
Retry Limit   : 0                  Retry Timer    : 30 sec
Signaling     : RSVP              Resv. Style    : FF
Hop Limit     : 255               Adaptive       : Enabled
FastReroute   : Disabled           Oper FR        : Disabled
FR Method     : Facility          FR Hop Limit   : 45
FR Bandwidth  : 0 Mbps            FR Node Protect: Disabled
FR Object     : Enabled
CSPF          : Disabled
Metric        : 1
Include Grps  :
None
ADSPEC        : Disabled
Use TE metric : Disabled
Exclude Grps  :
None

P2MPInstance:Test_p2mp                p2mp-inst-type : primary

S2L Name      :Test-s2l1             To              : 10.20.1.6
S2L Name      :Test-s2l2             To              : 10.20.1.5
S2L Name      :Test-s2l3             To              : 10.20.1.4
-----

A:ALU-25#

A:ALU-25# show router mpls p2mp-lsp Test_p2mp
=====
MPLS P2MP Instance (Originating)
=====
Type : Originating

```

```

-----
LSP Name       : lsp_1                      LSP Tunnel ID  : 1
P2MP ID        : 18                         Path LSP ID    : 18
Adm State      : Up                         Oper State     : Down

P2MPinstance:Test_p2mp                     p2mp-inst-type : primary
Inst Name      : lsp_1                     P2MP Inst ID   : 1
Adm State      : Up                         Oper State     : Down
Inst Up Time   : 0d 00:00:00               Inst Down Time : 0d 20:39:48
Hop Limit      : 255                       Adaptive       : Enabled
Record Route   : Record                    Record Label   : Record
Include Grps   :                           Exclude Grps   :
None                                                    None
Bandwidth      : 0 Mbps                     Oper Bw        : 0 Mbps

S2L Name       :Test-s2l1                   To              : 10.20.1.6
S2L Name       :Test-s2l2                   To              : 10.20.1.5
S2L Name       :Test-s2l3                   To              : 10.20.1.4
-----
A:ALU-25#

```

Note that the normal output is in detailed format only. There is no separate detail format.

A:ALU-52# show router mpls p2mp-lsp [p2mp-lsp-name] p2mp-instance [p2mp-inst-name]

=====

MPLS P2MP Instance (Originating)

=====

Type : Originating

```

-----
LSP Name       : lsp_1                      LSP Tunnel ID  : 1
P2MP ID        : 18                         Path LSP ID    : 18
Adm State      : Up                         Oper State     : Down

P2MPinstance:Test_p2mp                     p2mp-inst-type : primary
Inst Name      : lsp_1                     P2MP Inst ID   : 1
Adm State      : Up                         Oper State     : Down
Inst Up Time   : 0d 00:00:00               Inst Down Time : 0d 20:39:48
Hop Limit      : 255                       Adaptive       : Enabled
Record Route   : Record                    Record Label   : Record
Include Grps   :                           Exclude Grps   :
None                                                    None
Bandwidth      : 0 Mbps                     Oper Bw        : 0 Mbps

S2L Name       :Test-s2l1                   To              : 10.20.1.6
S2L Name       :Test-s2l2                   To              : 10.20.1.5
S2L Name       :Test-s2l3                   To              : 10.20.1.4
-----

```

A:ALU-52#

A:ALU-52# show router mpls p2mp-lsp [p2mp-lsp-name] p2mp-instance [p2mp-inst-name] mbb

=====

MPLS P2MP Instance (Originating)

=====

Type : Originating

=====

```

LSP Name       : lsp_1
P2MP ID        : 18
Adm State      : Up

LSP Tunnel ID  : 1
Path LSP ID    : 18
Oper State     : Down

P2MPinstance:Test_p2mp
Inst Name      : lsp_1
Adm State      : Up
Inst Up Time   : 0d 00:00:00
Hop Limit      : 255
Record Route   : Record
Include Grps   :
None
Bandwidth      : 0 Mbps
Last MBB       :
MBB type       :
ended at       :
In Prog MBB    :
MBB type       :
Started at     :
Failure code    :

p2mp-inst-type : primary
P2MP Inst ID   : 1
Oper State     : Down
Inst Down Time : 0d 20:39:48
Adaptive       : Enabled
Record Label   : Record
Exclude Grps   :
None
Oper Bw        : 0 Mbps

Mbb State      :
Old Metric     :

Next Retry In  :
Retry Attempt  :
Failure Node   :

S2L Name       :Test-s2l1
S2l Admin      :
Failure code    :

To              : 10.20.1.6
S2l Oper       :
Failure Node    : 10.12.1.1

S2L Name       :Test-s2l1
S2l Admin      :
Failure code    :

To              : 10.20.1.6
S2l Oper       :
Failure Node    : 10.12.1.1

```

A:ALU-52#

A:ALU-52# show router mpls p2mp-lsp [p2mp-lsp-name] p2mp-instance [p2mp-inst-name] s2l
[s2l-name]

=====

MPLS S2Ls (Originating)

```

=====
S2L Name          To          Next Hop      Adm   Opr
-----
Test-s2l1         10.20.1.6    10.10.1.2    Up    Up
-----

```

LSPs : 1

=====

A:ALU-52#

A:ALU-52# show router mpls p2mp-lsp [p2mp-lsp-name] p2mp-instance [p2mp-inst-name] s2l
[s2l-name] detail

=====

MPLS S2Ls (Originating) (Detail)

Type : Originating

```

-----
LSP Name       : lsp_1
P2MP ID        : 18
Adm State      : Up

LSP Tunnel ID  : 1
Path LSP ID    : 18
Oper State     : Down

```

```

P2MP Primary Instance:
Inst Name      : lsp_1
Adm State      : Up
P2MP Inst ID   : 1
Oper State     : Down

S2L Name       : Test-s2l1
Adm State      : Up
OutInterface   : 1/1/1
S2L Up Time    : 0d 00:00:00
Transitions    : 0
Retry Limit    : 0
RetryAttempt   : 0
Bandwidth      : No Reservation
Hop Limit      : 255
Record Route   : Record
Oper MTU       : 1496
FastReroute    : Disabled
FR Method      : Facility
FR Bandwidth   : 0 Mbps
FR Object      : Enabled
CSPF           : Disabled
Metric         : 1
Include Grps   :
None
CSPF Queries   : 9
Failure Code    : noError
ExplicitHops    :
    No Hops Specified
Actual Hops     :
    10.10.1.1(10.20.1.1) @
    -> 10.10.1.2(10.20.1.2)
ComputedHops    :
    10.10.1.1      -> 10.10.1.2
LastResignal    : n/a
P2MP Inst ID    : 1
Oper State      : Down
To              : 10.20.1.6
Oper State      : Down
Out Label       : 131071
S2L Down Time   : 0d 20:39:48
Path Changes    : 0
Retry Timer     : 30 sec
NextRetryIn     : 0 sec
Oper Bw         : 0 Mbps
Adaptive        : Enabled
Record Label    : Record
Neg MTU         : 1496
Oper FR         : Disabled
FR Hop Limit    : 45
FR Node Protect : Disabled
ADSPEC          : Disabled
Use TE metric   : Disabled
Exclude Grps    :
None
Failure Node    : n/a
Record Label    : N/A
Record Label    : 131071
CSPF Metric     : 1000
-----
A:ALU-52#

```

*A:Dut-C# show router mpls p2mp-lsp "ipmsi-1-73752" detail

```

=====
MPLS P2MP LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name       : ipmsi-1-73752
LSP Type       : P2mpAutoLsp
From           : 10.20.1.3
Adm State      : Up
LSP Up Time    : 0d 00:00:51
Transitions    : 3
Retry Limit    : 0
Signaling      : RSVP
Hop Limit      : 255
Adaptive       : Enabled
FastReroute    : Enabled
FR Method      : Facility
FR Node Pro*   : Disabled
LSP Tunnel ID  : 61445
Oper State     : Up
LSP Down Time  : 0d 00:00:00
Path Changes   : 3
Retry Timer    : 30 sec
Resv. Style    : SE
Negotiated MTU : n/a
ClassType      : 0
Oper FR        : Enabled
FR Hop Limit   : 16
FR Prop Adm Grp : Disabled

```

```

FR Object      : Enabled
Egress Stats: Enabled
CSPF           : Enabled
Metric         : Disabled
Include Grps:
None
Least Fill    : Disabled

Auto BW        : Disabled
LdpOverRsvp   : Enabled
IGP Shortcut: Enabled
IGP LFA       : Disabled
BGPTransTun   : Enabled
Oper Metric   : Disabled
Prop Adm Grp  : Disabled

Egress Oper St*: Out-of-resource
ADSPEC        : Disabled
Use TE metric  : Disabled
Exclude Grps   :
None

VprnAutoBind  : Enabled
BGP Shortcut   : Enabled
IGP Rel Metric : Disabled

P2MPInstance: 1
S2L Cfg Cou*: 4
S2l-Name      : path_ipmsi
S2l-Name      : path_ipmsi
S2l-Name      : path_ipmsi
S2l-Name      : path_ipmsi

P2MP-Inst-type : Primary
S2L Oper Count*: 4
To              : 10.20.1.1
To              : 10.20.1.2
To              : 10.20.1.5
To              : 10.20.1.6
=====

```

sr-te-lsp

Syntax **sr-te-lsp** [*lsp-name*] [**detail**]
sr-te-lsp [*lsp-name*] [**status** {**up** | **down**}] [**detail**] **path** [*path-name*]
sr-te-lsp [*lsp-name*] [**status** {**up** | **down**}] [**to** *ip-address*] [**detail**]

Context show>router>mpls

Description This command displays SR-TE-LSP information.

Parameters **lsp-name** — Specifies the name of the SR-TE LSP, up to 64 characters.
detail — Displays detailed information.
status — Allows filtering of the output based on the status of the SR-TE LSP.

Values **up**, **down**

path *path-name* — Displays SR-TE LSP path information.

to *ip-address* — Displays the IP address of the egress router for the LSP.

Output

Sample Output

The following output applies to router- or PCC-controlled SR-TE LSPs.

```

*A:Dut-A# show router mpls sr-te-lsp "srTeRtrControlled" path "1" detail
=====
MPLS SR-TE LSP srTeRtrControlled Path 1 (Detail)
=====

```

```

Legend :
    S - Strict                      L - Loose
=====
-----
SR-TE LSP srTeRtrControlled Path 1
-----
LSP Name       : srTeRtrControlled
Path LSP ID    : 1024
From           : 10.20.1.1          To           : 10.20.1.6
Admin State    : Up                Oper State   : Up
Path Name      : 1                 Path Type    : Primary
Path Admin     : Up                Path Oper    : Up
Path Up Time   : 0d 00:03:19       Path Down Time : 0d 00:00:00
Retry Limit    : 0                 Retry Timer   : 30 sec
Retry Attempt  : 1                 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

Include Groups :                    Oper Include Groups :
None                                           None
Exclude Groups :                    Oper Exclude Groups :
None                                           None

IGP/TE Metric  : 16777215          Oper Metric     : 16777215
Oper MTU       : 1484              Path Trans      : 1
Failure Code    : noError
Failure Node    : n/a
Explicit Hops   :
    10.20.1.2(S)      -> 10.20.1.3(L)
Actual Hops     :
    10.10.1.2 (10.20.1.2)          Record Label   : 262143
    -> 10.20.1.3 (10.20.1.3)       Record Label   : 102003
    -> 10.20.1.6 (10.20.1.6)       Record Label   : 103006
=====

*A:Dut-A# show router mpls sr-te-lsp "srTeRtrControlled" detail
=====
MPLS SR-TE LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name       : srTeRtrControlled
LSP Type       : SrTeLsp           LSP Tunnel ID  : 2
LSP Index      : 65537             TTM Tunnel Id  : 655363
From           : 10.20.1.1          To             : 10.20.1.6

```



```

Adm State      : Up
LSP Up Time    : 0d 00:06:09
Transitions    : 1
Retry Limit    : 0
Hop Limit      : 255
CSPF           : Disabled
Metric         : N/A
Include Grps   :
None

Oper State     : Up
LSP Down Time  : 0d 00:00:00
Path Changes   : 1
Retry Timer    : 30 sec
Negotiated MTU : 1484

Exclude Grps   :
None

VprnAutoBind  : Enabled
IGP Shortcut   : Enabled
IGP LFA        : Disabled
BGPTransTun   : Enabled
Oper Metric    : 16777215
PCE Report     : Inherited
PCE Compute    : Disabled
Path Profile   :
None

BGP Shortcut   : Enabled
IGP Rel Metric : Disabled

Max SR Labels  : 6
PCE Control    : Disabled

Primary(a)     : 1
Bandwidth      : 0 Mbps
Up Time        : 0d 00:06:09

=====
Outputs with other filters:
-----
*A:Dut-A# show router mpls sr-te-lsp
=====
MPLS SR-TE LSPs (Originating)
=====
LSP Name      To      Tun    Protect  Adm  Opr
                Id      Path
-----
sr-tel        10.20.1.6    1      N/A      Up   Up
srTeRtrControlled 10.20.1.6    2      N/A      Up   Up
-----
LSPs : 2
=====

*A:Dut-A# show router mpls sr-te-lsp path
=====
MPLS SR-TE LSP Path
=====
LSP Name      : sr-tel      To      : 10.20.1.6
Adm State     : Up          Oper State : Up
-----
Path Name      Type      Adm  Opr
-----
1              Primary    Up   Up
-----
LSP Name      : srTeRtrControlled To      : 10.20.1.6
Adm State     : Up          Oper State : Up
-----
Path Name      Type      Adm  Opr
-----
1              Primary    Up   Up
-----

*A:Dut-A# show router mpls sr-te-lsp to 10.20.1.6

```

```

=====
MPLS SRTE LSPs (Originating)
=====
LSP Name                               To                               Tun    Protect  Adm  Opr
                                   Id    Path
-----
sr-te1                                10.20.1.6                        1      N/A      Up   Up
srTeRtrControlled                     10.20.1.6                        2      N/A      Up   Up
-----
LSPs : 2
=====
*A:Dut-A# show router mpls sr-te-lsp status up
=====
MPLS SR-TE LSPs (Originating)
=====
LSP Name                               To                               Tun    Protect  Adm  Opr
                                   Id    Path
-----
sr-te1                                10.20.1.6                        1      N/A      Up   Up
srTeRtrControlled                     10.20.1.6                        2      N/A      Up   Up
-----
LSPs : 2
=====
*A:Dut-A# show router mpls sr-te-lsp "srTeRtrControlled" status up detail path "1"
=====
MPLS SR-TE LSP srTeRtrControlled Path 1 (Detail)
=====
Legend :
    S - Strict                      L - Loose
=====
SR-TE LSP srTeRtrControlled Path 1
-----
LSP Name          : srTeRtrControlled
Path LSP ID       : 1024
From              : 10.20.1.1          To              : 10.20.1.6
Admin State       : Up                Oper State       : Up
Path Name         : 1                 Path Type        : Primary
Path Admin        : Up                Path Oper        : Up
Path Up Time      : 0d 00:07:12        Path Down Time   : 0d 00:00:00
Retry Limit       : 0                 Retry Timer      : 30 sec
Retry Attempt     : 1                 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

```

Include Groups :	None	Oper Include Groups :	None
Exclude Groups :	None	Oper Exclude Groups :	None
IGP/TE Metric :	16777215	Oper Metric :	16777215
Oper MTU :	1484	Path Trans :	1
Failure Code :	noError		
Failure Node :	n/a		
Explicit Hops :			
10.20.1.2(S)	-> 10.20.1.3(L)		
Actual Hops :			
10.10.1.2 (10.20.1.2)		Record Label :	262143
-> 10.20.1.3 (10.20.1.3)		Record Label :	102003
-> 10.20.1.6 (10.20.1.6)		Record Label :	103006

=====

The following output applies to PCE-computed SR-TE LSPs.

*A:Dut-B>config>router>mpls>lsp# /show router mpls sr-te-lsp "srte_lsp" path detail

=====

MPLS SR-TE LSP srte_lsp Path (Detail)

=====

Legend :

S - Strict

L - Loose

SR-TE LSP srte_lsp Path fully_loose

LSP Name	: srte_lsp		
Path LSP ID	: 2		
From	: 10.20.1.2	To	: 10.20.1.5
Admin State	: Up	Oper State	: Up
Path Name	: fully_loose	Path Type	: Primary
Path Admin	: Up	Path Oper	: Up
Path Up Time	: 0d 10:48:05	Path Down Time	: 0d 00:00:00
Retry Limit	: 0	Retry Timer	: 30 sec
Retry Attempt	: 1	Next Retry In	: 0 sec
CSPF	: Enabled	Oper CSPF	: Enabled
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	: Enabled	Oper PCE Report	: Enabled
PCE Control	: Disabled	Oper PCE Control	: Disabled
PCE Compute	: Enabled		

Include Groups :	None	Oper Include Groups :	None
Exclude Groups :	None	Oper Exclude Groups :	None
IGP/TE Metric :	20	Oper Metric :	20

```

Oper MTU          : 1492                      Path Trans          : 1
Failure Code      : noError
Failure Node      : n/a
Explicit Hops     :
    No Hops Specified
Actual Hops       :
    1.1.2.1                      Record Label          : 131071
-> 1.1.5.5          Record Label          : 131068
=====

*A:Dut-B>config>router>mpls>lsp# /show router mpls sr-te-lsp "srte_lsp" detail
=====
MPLS SR-TE LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name          : srte_lsp
LSP Type          : SrTeLsp                    LSP Tunnel ID       : 1
LSP Index         : 65536                      TTM Tunnel Id      : 655362
From              : 10.20.1.2                  To                 : 10.20.1.5
Adm State         : Up                        Oper State          : Up
LSP Up Time       : 0d 10:48:17                LSP Down Time      : 0d 00:00:00
Transitions       : 1                        Path Changes        : 1
Retry Limit       : 0                        Retry Timer         : 30 sec
Hop Limit         : 255                      Negotiated MTU      : 1492
CSPF              : Enabled
Metric            : N/A                      Use TE metric       : Disabled
Include Grps      :                          Exclude Grps        :
None                                     None
VprnAutoBind      : Enabled
IGP Shortcut      : Enabled                    BGP Shortcut        : Enabled
IGP LFA           : Disabled                    IGP Rel Metric      : Disabled
BGPTransTun       : Enabled
Oper Metric       : 20
PCE Report        : Enabled
PCE Compute       : Enabled                    PCE Control         : Disabled
Max SR Labels     : 6                        Additional FRR Labels: 1
Path Profile      :
None
Primary(a)        : fully_loose                Up Time             : 0d 10:48:17
Bandwidth         : 0 Mbps
=====

```

The following output applies to a PCE-computed SR-TE LSP when **pce-report sr-te** is enabled under MPLS.

```

*A:Dut-B>config>router>mpls>lsp# /show router mpls sr-te-lsp "srte_lsp" path detail
=====
MPLS SR-TE LSP srte_lsp Path (Detail)
=====
Legend :
    S - Strict                      L - Loose
=====
-----
SR-TE LSP srte_lsp Path fully_loose
-----
LSP Name          : srte_lsp
Path LSP ID       : 2

```

```

From          : 10.20.1.2          To          : 10.20.1.5
Admin State   : Up                 Oper State   : Up
Path Name     : fully_loose        Path Type    : Primary
Path Admin    : Up                 Path Oper    : Up
Path Up Time  : 0d 10:51:47        Path Down Time : 0d 00:00:00
Retry Limit   : 0                  Retry Timer   : 30 sec
Retry Attempt : 1                  Next Retry In : 0 sec
CSPF          : Enabled            Oper CSPF     : Enabled
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          Oper PCE Report : Enabled
PCE Report    : Inherited          Oper PCE Control : Disabled
PCE Control   : Disabled
PCE Compute   : Enabled

Include Groups :                    Oper Include Groups :
None           :                    None
Exclude Groups :                    Oper Exclude Groups :
None           :                    None

IGP/TE Metric : 20                  Oper Metric      : 20
Oper MTU       : 1492               Path Trans       : 1
Failure Code   : noError
Failure Node   : n/a
Explicit Hops  :
    No Hops Specified
Actual Hops    :
    1.1.2.1
-> 1.1.5.5
Record Label   : 131071
Record Label   : 131068
=====

*A:Dut-B>config>router>mpls>lsp# /show router mpls sr-te-lsp "srte_lsp" detail
=====
MPLS SR-TE LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
-----
LSP Name      : srte_lsp
LSP Type      : SrTeLsp
LSP Index     : 65536
LSP Tunnel ID : 1
LSP Tunnel Id : 655362
From          : 10.20.1.2
To            : 10.20.1.5
Adm State     : Up
Oper State     : Up
LSP Up Time   : 0d 10:52:16
LSP Down Time : 0d 00:00:00
Transitions   : 1
Path Changes   : 1
Retry Limit   : 0
Retry Timer    : 30 sec
Hop Limit     : 255
Negotiated MTU : 1492
CSPF          : Enabled
Metric        : N/A
Use TE metric  : Disabled
Include Grps   :
Exclude Grps   :
None           :
VprnAutoBind  : Enabled
IGP Shortcut   : Enabled
BGP Shortcut   : Enabled
IGP LFA        : Disabled
IGP Rel Metric : Disabled

```

```

BGPTransTun      : Enabled
Oper Metric      : 20
PCE Report       : Inherited
PCE Compute      : Enabled
Max SR Labels    : 6
Path Profile     :
None
Primary(a)       : fully_loose
Bandwidth        : 0 Mbps
PCE Control      : Disabled
Additional FRR Labels: 1
Up Time          : 0d 10:52:16

```

The following output applies to PCE-controlled SR-TE LSPs.

```

*A:Dut-B>config>router>pcep>pcc# /show router mpls sr-te-lsp "srte_lsp" path detail
=====
MPLS SR-TE LSP srte_lsp Path (Detail)
=====
Legend :
    S - Strict          L - Loose
=====
SR-TE LSP srte_lsp Path fully_loose
-----
LSP Name          : srte_lsp
Path LSP ID       : 2
From              : 10.20.1.2
Admin State       : Up
Path Name         : fully_loose
Path Admin        : Up
Path Up Time      : 0d 00:01:07
Retry Limit       : 0
Retry Attempt     : 1
CSPF              : Enabled
Bandwidth         : No Reservation
Hop Limit         : 255
Setup Priority    : 7
Hold Priority     : 0
Inter-area        : N/A

To                : 10.20.1.5
Oper State        : Up
Path Type         : Primary
Path Oper         : Up
Path Down Time    : 0d 00:00:00
Retry Timer       : 30 sec
Next Retry In     : 0 sec
Oper CSPF         : Enabled
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        : Enabled
PCE Control       : Enabled
PCE Compute       : Enabled

PCE Updt State    : None
Oper PCE Report   : Enabled
Oper PCE Control  : Enabled

Include Groups    :
None
Exclude Groups   :
None

Oper Include Groups :
None
Oper Exclude Groups :
None

IGP/TE Metric     : 20
Oper MTU          : 1492
Failure Code      : noError
Failure Node      : n/a
Explicit Hops     :
    No Hops Specified
Actual Hops       :
    1.1.2.1
-> 1.1.5.5

Oper Metric       : 20
Path Trans        : 1

Record Label      : 131071
Record Label      : 131068

```

```

=====
*A:Dut-B>config>router>pcep>pcc# /show router mpls sr-te-lsp "srte_lsp" detail
=====
MPLS SR-TE LSPs (Originating) (Detail)
=====
-----
Type : Originating
-----
LSP Name       : srte_lsp
LSP Type       : SrTeLsp
LSP Index      : 65536
From           : 10.20.1.2
Adm State      : Up
LSP Up Time    : 0d 00:01:38
Transitions    : 1
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    : 20
PCE Report     : Enabled
PCE Compute    : Enabled
Max SR Labels  : 6
Path Profile   :
None
Primary(a)    : fully_loose
Bandwidth      : 0 Mbps

LSP Tunnel ID  : 1
TTM Tunnel Id  : 655362
To             : 10.20.1.5
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

BGP Shortcut   : Enabled
IGP Rel Metric : Disabled

PCE Control    : Enabled
Additional FRR Labels: 1

Up Time        : 0d 00:01:38
=====

*A:Dut-B>config>router>mpls>lsp# /show router pcep pcc lsp-db
=====
PCEP Path Computation Client (PCC) LSP Update Info
=====
PCEP-specific LSP ID: 3
LSP ID         : 2
Tunnel ID      : 1
LSP Name       : srte_lsp::fully_loose
Source Address : 10.20.1.2
LSP Delegated  : True
Oper Status    : active

LSP Type       : seg-rt
Extended Tunnel Id : 10.20.1.2
Destination Address : 10.20.1.5
Delegate PCE Address: 10.20.1.24
=====

*A:Dut-B>config>router>mpls>lsp# /show router mpls sr-te-lsp "srte_lsp" path
=====
MPLS SR-TE LSP srte_lsp Path
=====
-----
LSP Name       : srte_lsp
Adm State      : Up

To             : 10.20.1.5
Oper State     : Up

-----
Path Name      Type      Adm  Opr
-----
fully_loose    Primary    Up   Up

```

=====

srkg-database

- Syntax** **srkg-database** [**router-id** *ip-address*] [**interface** *ip-address*]
- Context** show>router>mpls
- Description** This command displays MPLS SRLG database information.
- Parameters** **router-id** *ip-address* — Specifies a 32-bit integer uniquely identifying the router in the Autonomous System. By convention to ensure uniqueness, this may default to the value of one of the router's IPv4 host addresses, represented as a 32-bit unsigned integer, if IPv4 is configured on the router. The router-id can be either the local one or some remote router.
- interface** *ip-address* — Specifies the IP address of the interface.

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* — Keyword to display binding information.
- Output** MPLS Path Output

[Table 21](#) describes MPLS Path output fields.

Table 21 MPLS Path Output Fields

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.

Table 21 MPLS Path Output Fields (Continued)

Label	Description (Continued)
Strict/Loose	Strict — The LSP must take a direct path from the previous hop router to the next router. Loose — The route taken by the LSP from the previous hop to the next hop can traverse through other routers.
LSP Name	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.

Sample Output

```
*A:SRU4>config>router>mpls# show router mpls path
=====
MPLS Path:
=====
Path Name                               Adm  Hop Index  IP Address  Strict/Loose
-----
to_110_20_1_1                          Up    no hops    n/a         n/a
to_110_20_1_2                          Up    no hops    n/a         n/a
to_110_20_1_3                          Up    no hops    n/a         n/a
to_110_20_1_4                          Up    no hops    n/a         n/a
to_110_20_1_5                          Up    no hops    n/a         n/a
to_110_20_1_6                          Up    no hops    n/a         n/a
to_110_20_1_110                        Up    no hops    n/a         n/a
to_10_8_100_15                         Up    no hops    n/a         n/a
to_10_20_1_20                          Up    no hops    n/a         n/a
to_10_20_1_22                          Up    no hops    n/a         n/a
to_10_100_1_1                          Up    no hops    n/a         n/a
-----
Paths : 11
=====
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls path lsp-binding
=====
MPLS Path:
=====
Path Name                               Opr  LSP Name  Binding
-----
to_110_20_1_1                          Up    to_110_20_1_1_cspf  Primary
to_110_20_1_1                          Up    to_110_20_1_1_cspf_2  Primary
to_110_20_1_1                          Up    to_110_20_1_1_cspf_3  Primary
Up    to_110_20_1_1_cspf_16              Primary
to_110_20_1_1                          Up    to_110_20_1_1_cspf_17  Primary
to_110_20_1_1                          Up    to_110_20_1_1_cspf_18  Primary
to_110_20_1_1                          Up    to_110_20_1_1_cspf_19  Primary
to_110_20_1_1                          Up    to_110_20_1_1_cspf_20  Primary
to_110_20_1_2                          Up    to_110_20_1_2_cspf    Primary
```

```

Up      to_110_20_1_2_cspf_2      Primary
Up      to_110_20_1_2_cspf_3      Primary
Up      to_110_20_1_2_cspf_4      Primary
Up      to_110_20_1_2_cspf_5      Primary
...
to_10_100_1_1      Down to_10_100_1_1_cspf      Primary
                        Down to_10_100_1_1_cspf_2      Primary
                        Down to_10_100_1_1_cspf_3      Primary
                        Down to_10_100_1_1_cspf_4      Primary
                        Down to_10_100_1_1_cspf_5      Primary
                        Down to_10_100_1_1_cspf_6      Primary
Down to_10_100_1_1_cspf_13      Primary
                        Down to_10_100_1_1_cspf_14      Primary
                        Down to_10_100_1_1_cspf_15      Primary
                        Down to_10_100_1_1_cspf_16      Primary
                        Down to_10_100_1_1_cspf_17      Primary
                        Down to_10_100_1_1_cspf_18      Primary
                        Down to_10_100_1_1_cspf_19      Primary
                        Down to_10_100_1_1_cspf_20      Primary
-----
Paths : 11
=====
*A:SRU4>config>router>mpls#

```

srlg-group

- Syntax** `srlg-group [group-name]`
- Context** `show>router>mpls`
- Description** This command displays MPLS SRLG groups
- Parameters** *group-name* — Specifies the name of the SRLG group within a virtual router instance.
- Output** MPLS SRLG Group Output

[Table 22](#) describes MPLS SRLG group output fields

Table 22 MPLS SRLG Group Output Fields

Label	Description
Group Name	Displays the name of the SRLG group within a virtual router instance.
Group Value	Displays the group value associated with this SRLG group.
Interface	Displays the interface where the SRLG groups is associated.
No. of Groups	Displays the total number of SRLG groups associated with the output.

Sample Output

```
*A:SRU4>config>router>mpls# show router mpls srlg-group
=====
MPLS Srlg Groups
=====
```

Group Name	Group Value	Interfaces
1432	1432	srl-1
1433	1433	srl-3
1434	1434	aps-8
1435	1435	aps-9
2410	2410	srr-1
2411	2411	srr-2
2412	2412	srr-3
3410	3410	aps-1
3420	3420	aps-2
3430	3430	aps-3
3440	3440	sr4-1
41.80	4180	g7600
41104	41104	germ-1
415.70	41570	gsr1
420.40	42040	m160
422.60	42260	gsr2
44.200	44200	hubA
45100	45100	ess-7-1
45110	45110	ess-7-2
45120	45120	ess-7-3
4651	4651	src-1.1
4652	4652	src-1.2
4653	4653	src-1.3
4654	4654	src-1.4
4655	4655	src-1.5
4656	4656	src-1.6
4657	4657	src-1.7
4658	4658	src-1.8
4659	4659	src-1.9
4660	4660	src-1.10

```
-----
No. of Groups: 30
=====
*A:SRU4>config>router>mpls#

*A:SRU4>config>router>mpls# show router mpls srlg-group "1432"
=====
MPLS Srlg Groups
=====
```

Group Name	Group Value	Interfaces
1432	1432	srl-1

```
-----
No. of Groups: 1
=====
*A:SRU4>config>router>mpls#
```

static-lsp

Syntax	static-lsp [lsp-name] static-lsp {transit terminate} static-lsp count
Context	show>router>mpls
Description	This command displays MPLS static LSP information.
Output	MPLS Static LSP Output

[Table 23](#) describes the MPLS static LSP output fields.

Table 23 MPLS Static LSP Output Fields

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.
In I/F	The ingress interface.
Out Label	The egress interface.
Out I/F	The egress interface.
Adm	Down — The path is administratively disabled. Up — The path is administratively enabled.
Opr	Down — The path is operationally down. Up — The path is operationally up.
LSPs	The total number of static LSPs.

Sample Output

```
A:ALA-12# show router mpls static-lsp
=====
MPLS Static LSPs (Originating)
=====
Lsp Name           To           Next Hop      Out Label  Out I/F    Adm  Opr
-----
NYC_SJC_customer2 100.20.1.10   10.10.1.4     1020      1/1/1      Up   Up
-----
LSPs : 1
=====
A:ALA-12#
```

```
*A:SRU4>config>router>mpls# show router mpls static-lsp transit
```

```
=====
MPLS Static LSPs (Transit)
=====
In Label    In Port    Out Label   Out Port    Next Hop      Adm  Opr
-----
240         aps-1      440         1/1/10      11.22.11.3    Up   Up
241         aps-1      441         1/1/10      11.22.11.3    Up   Up
242         aps-1      442         1/1/10      11.22.11.3    Up   Up
243         aps-1      443         1/1/10      11.22.11.3    Up   Up
244         aps-1      444         1/1/10      11.22.11.3    Up   Up
245         aps-1      445         1/1/10      11.22.11.3    Up   Up
246         aps-1      446         1/1/10      11.22.11.3    Up   Up
247         aps-1      447         1/1/10      11.22.11.3    Up   Up
248         aps-1      448         1/1/10      11.22.11.3    Up   Up
249         aps-1      449         1/1/10      11.22.11.3    Up   Up
250         aps-1      450         1/1/10      11.22.11.3    Up   Up
251         aps-1      451         1/1/10      11.22.11.3    Up   Up
252         aps-1      452         1/1/10      11.22.11.3    Up   Up
253         aps-1      453         1/1/10      11.22.11.3    Up   Up
...
207         3/2/8     407         1/1/9       11.22.10.3    Up   Up
208         3/2/8     408         1/1/9       11.22.10.3    Up   Up
209         3/2/8     409         1/1/9       11.22.10.3    Up   Up
-----
LSPs : 256
=====
*A:SRU4>config>router>mpls#

A:ALA-12# show router mpls static-lsp terminate
=====
MPLS Static LSPs (Terminate)
=====
In Label    In I/F     Out Label   Out I/F     Next Hop      Adm  Opr
-----
1021        1/1/1     n/a         n/a         n/a           Up   Up
-----
LSPs : 1
=====
A:ALA-12#
```

statistics-summary

Syntax	statistics-summary
Context	show>router>mpls>statistics-summary
Description	This command displays the number of LSP statistics configured.
Output	

Sample Output

```
*A:SRU4>config>router>mpls# show router mpls statistics-summary
=====
Statistics Summary
```

```

=====
LSP egress statistics           : 0
LSP ingress statistics         : 0
=====
*A:SRU4>config>router>mpls#

```

status

- Syntax** **status**
- Context** show>router>mpls
- Description** This command displays MPLS operation information.
- Output** MPLS Status Output

[Table 24](#) describes MPLS status output fields.

Table 24 **MPLS Status Output Fields**

Label	Description
Admin Status	Down — MPLS is administratively disabled. Up — MPLS is administratively enabled.
Oper Status	Down — MPLS is operationally down. Up — MPLS is operationally up.
LSP Counts	Static LSPs — Displays the count of static LSPs that originate, transit, and terminate on or through the router. Dynamic LSPs — Displays the count of dynamic LSPs that originate, transit, and terminate on or through the router. Detour LSPs — Displays the count of detour LSPs that originate, transit, and terminate on or through the router.
FR Object	Enabled — Specifies that Fast reroute object is signaled for the LSP. Disabled — Specifies that Fast reroute object is not signaled for the LSP.
Resignal Timer	Enabled — Specifies that the resignal timer is enabled for the LSP. Disabled — Specifies that the resignal timer is disabled for the LSP.
Hold Timer	Displays the amount of time that the ingress node holds before programming its data plane and declaring the LSP up to the service module.

Sample Output

```
*A:DutC>show router mpls status
```

```
=====
MPLS Status
=====
Admin Status      : Down          Oper Status      : Down
Oper Down Reason  : adminDown
FRR Object        : Enabled       Resignal Timer    : Disabled
Hold Timer        : 1 seconds     Next Resignal     : N/A
Srlg Frr          : Disabled      Srlg Frr Strict   : Disabled
Admin Group Frr   : Disabled
Dynamic Bypass    : Enabled       User Srlg Database : Disabled
BypassResignalTimer: Disabled    BypassNextResignal : N/A
Least Fill Min Thd.: 5 percent   LeastFill ReoptiThd: 10 percent
Local TTL Prop    : Enabled       Transit TTL Prop   : Enabled
AB Sample Multipli*: 1           AB Adjust Multipli*: 288
Exp Backoff Retry : Disabled      CSPF On Loose Hop  : Disabled
Lsp Init RetryTime*: 30 seconds   MBB Pref Current H*: Disabled
Logger Event Bundl*: Disabled
RetryIgpOverload  : Disabled

P2mp Resignal Timer: Disabled      P2mp Next Resignal : N/A
Sec FastRetryTimer : Disabled      Static LSP FR Timer: 30 seconds
P2P Max Bypass Ass*: 1000
P2PActPathFastRetry: Disabled      P2MP S2L Fast Retry: Disabled
In Maintenance Mode: No

MplsTp           : Disabled
Pce-report        : None
Next Available Lsp Index : 1
Entropy Label Rsvp-TE : Enabled

=====
MPLS LSP Counts
=====
```

	Originate	Transit	Terminate
Static LSPs	0	0	0
Dynamic LSPs	1	0	1
Detour LSPs	0	0	0
P2MP S2Ls	0	0	0
MPLS-TP LSPs	0	0	0
Mesh-P2P LSPs	0	N/A	N/A
One Hope-P2P LSPs	0	N/A	N/A
SR-TE LSPs	0	N/A	N/A

```
=====
```

tp-lsp

Syntax	tp-lsp [<i>lsp-name</i>] [status { up down }] [from <i>ip-address</i> to <i>ip-address</i>] [detail] tp-lsp [<i>lsp-name</i>] path [protecting working] [detail] tp-lsp [<i>lsp-name</i>] protection
Context	show>router>mpls
Description	This command displays TP LSP information.

- Parameters**
- lsp-name* — Specifies the LSP name up to 32 characters; accepts * as a wild card character
 - path* — Displays LSP path information.
 - protection* — Displays LSP protection information.
 - up | down* — Specifies the state of the LSP.

Output

Sample Output

```
*A:mlstp-dutA# show router mpls tp-lsp
path
protection
to <a.b.c.d>
<lsp-name>
  "lsp-32" "lsp-33" "lsp-34" "lsp-35" "lsp-36" "lsp-37" "lsp-38" "lsp-39"
  "lsp-40" "lsp-41"
status {up | down}
from <ip-address>
detail

*A:mlstp-dutA# show router mpls tp-lsp "lsp-
"lsp-32" "lsp-33" "lsp-34" "lsp-35" "lsp-36" "lsp-37" "lsp-38" "lsp-39"
"lsp-40" "lsp-41"
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32"

=====
MPLS MPLS-TP LSPs (Originating)
=====
LSP Name                               To                Tun      Protect   Adm  Opr
                                Id                Id      Path
-----
lsp-32                                0.0.3.234         32       No        Up   Up
-----
LSPs : 1
=====
*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" detail

=====
MPLS MPLS-TP LSPs (Originating) (Detail)
=====
Type : Originating
-----
LSP Name      : lsp-32
LSP Type      : MplsTp
From Node Id: 0.0.3.233+
Adm State     : Up
LSP Up Time   : 0d 04:50:47
Transitions   : 1
DestGlobalId: 42
LSP Tunnel ID : 32
To Node Id    : 0.0.3.234
Oper State    : Up
LSP Down Time : 0d 00:00:00
Path Changes   : 2
DestTunnelNum : 32

=====
*A:mlstp-dutA# show router mpls tp-lsp path
```



```

=====
MPLS-TP LSP Path Information
=====
LSP Name       : lsp-32                               To       : 0.0.3.234
Admin State    : Up                                   Oper State : Up

-----
Path           NextHop           InLabel   OutLabel   Out I/F           Admin   Oper
-----
Working                32         32         AtoB_1         Up     Down
Protect               2080        2080        AtoC_1         Up      Up
=====
LSP Name       : lsp-33                               To       : 0.0.3.234
Admin State    : Up                                   Oper State : Up

-----
Path           NextHop           InLabel   OutLabel   Out I/F           Admin   Oper
-----
Working                33         33         AtoB_1         Up     Down
Protect               2082        2082        AtoC_1         Up      Up
=====
LSP Name       : lsp-34                               To       : 0.0.3.234
Admin State    : Up                                   Oper State : Up

-----
Path           NextHop           InLabel   OutLabel   Out I/F           Admin   Oper
-----
Working                34         34         AtoB_1         Up     Down
Protect               2084        2084        AtoC_1         Up      Up
=====
LSP Name       : lsp-35                               To       : 0.0.3.234
Admin State    : Up                                   Oper State : Up

-----
Path           NextHop           InLabel   OutLabel   Out I/F           Admin   Oper
-----
Working                35         35         AtoB_1         Up     Down
Protect               2086        2086        AtoC_1         Up      Up
=====
LSP Name       : lsp-36                               To       : 0.0.3.234
Admin State    : Up                                   Oper State : Up

-----
Path           NextHop           InLabel   OutLabel   Out I/F           Admin   Oper
-----
Working                36         36         AtoB_1         Up     Down
Protect               2088        2088        AtoC_1         Up      Up
=====
LSP Name       : lsp-37                               To       : 0.0.3.234
Admin State    : Up                                   Oper State : Up

-----
Path           NextHop           InLabel   OutLabel   Out I/F           Admin   Oper
-----
Working                37         37         AtoB_1         Up     Down
Protect               2090        2090        AtoC_1         Up      Up
=====
LSP Name       : lsp-38                               To       : 0.0.3.234

```

Admin State : Up Oper State : Up

Path	NextHop	InLabel	OutLabel	Out I/F	Admin	Oper
Working		38	38	AtoB_1	Up	Down
Protect		2092	2092	AtoC_1	Up	Up

=====

LSP Name : lsp-39 To : 0.0.3.234

Admin State : Up Oper State : Up

Path	NextHop	InLabel	OutLabel	Out I/F	Admin	Oper
Working		39	39	AtoB_1	Up	Down
Protect		2094	2094	AtoC_1	Up	Up

=====

LSP Name : lsp-40 To : 0.0.3.234

Admin State : Up Oper State : Up

Path	NextHop	InLabel	OutLabel	Out I/F	Admin	Oper
Working		40	40	AtoB_1	Up	Down
Protect		2096	2096	AtoC_1	Up	Up

=====

LSP Name : lsp-41 To : 0.0.3.234

Admin State : Up Oper State : Up

Path	NextHop	InLabel	OutLabel	Out I/F	Admin	Oper
Working		41	41	AtoB_1	Up	Down
Protect		2098	2098	AtoC_1	Up	Up

*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path working

=====

MPLS-TP LSP Working Path Information

LSP: "lsp-32"

=====

LSP Name : lsp-32 To : 0.0.3.234

Admin State : Up Oper State : Up

Path	NextHop	InLabel	OutLabel	Out I/F	Admin	Oper
Working		32	32	AtoB_1	Up	Down

=====

*A:mlstp-dutA# show router mpls tp-lsp "lsp-32" path protect

=====

MPLS-TP LSP Protect Path Information

LSP: "lsp-32"

=====

LSP Name : lsp-32 To : 0.0.3.234

Admin State : Up Oper State : Up

=====

```
Path          NextHop          InLabel  OutLabel  Out I/F          Admin  Oper
-----
Protect                2080      2080      AtoC_1          Up     Up
=====
*A:mlstp-dutA#  show router mpls tp-lsp "lsp-32" path protect detail

=====
MPLS-TP LSP Protect Path Information
      LSP: "lsp-32" (Detail)
=====
LSP Name       : lsp-32                      To           : 0.0.3.234
Admin State    : Up                          Oper State    : Up

Protect path information
-----
Path Type      : Protect                      LSP Num      : 2
Path Admin     : Up                          Path Oper     : Up
Out Interface  : AtoC_1                      Next Hop Addr : n/a
In Label       : 2080                        Out Label     : 2080
Path Up Time   : 0d 04:52:17                 Path Dn Time  : 0d 00:00:00
Active Path    : Yes                          Active Time   : 0d 00:52:56

MEP information
MEP State      : Up                          BFD           : cc
OAM Templ      : privatebed-oam-template      CC Status     : inService
                                           CV Status     : unknown
Protect Templ  : privatebed-protection-template WTR Count Down: 0 seconds
RX PDU         : SF (1,1)                     TX PDU        : SF (1,1)
Defects        :
=====
*A:mlstp-dutA#  show router mpls tp-lsp "lsp-32" path working detail

=====
MPLS-TP LSP Working Path Information
      LSP: "lsp-32" (Detail)
=====
LSP Name       : lsp-32                      To           : 0.0.3.234
Admin State    : Up                          Oper State    : Up

Working path information
-----
Path Type      : Working                      LSP Num      : 1
Path Admin     : Up                          Path Oper     : Up
Out Interface  : AtoB_1                      Next Hop Addr : n/a
In Label       : 32                          Out Label     : 32
Path Up Time   : 18d 22:39:53                 Path Dn Time  : 0d 00:00:00
Active Path    : Yes                          Active Time   : 18d 22:34:52

MEP information
MEP State      : Up                          BFD           : cc
OAM Templ      : privatebed-oam-template      BFD Status    : inService
Ingress I/F N*: 0                          Egress I/F Num: 0
AIS State      : AIS
=====
*A:mlstp-dutA#

*A:mlstp-dutA#  show router mpls tp-lsp protection
```

```
=====
MPLS-TP LSP Protection Information
Legend: W-Working, P-Protect,
=====
```

LSP Name	Admin State	Oper State	Path State	Ingr/Egr Label	Act. Path	Rx Tx	PDU PDU

lsp-32	Up	Up	W Down	32/32	No	SF	(1,1)
			P Up	2080/2080	Yes	SF	(1,1)
lsp-33	Up	Up	W Down	33/33	No	SF	(1,1)
			P Up	2082/2082	Yes	SF	(1,1)
lsp-34	Up	Up	W Down	34/34	No	SF	(1,1)
			P Up	2084/2084	Yes	SF	(1,1)
lsp-35	Up	Up	W Down	35/35	No	SF	(1,1)
			P Up	2086/2086	Yes	SF	(1,1)
lsp-36	Up	Up	W Down	36/36	No	SF	(1,1)
			P Up	2088/2088	Yes	SF	(1,1)
lsp-37	Up	Up	W Down	37/37	No	SF	(1,1)
			P Up	2090/2090	Yes	SF	(1,1)
lsp-38	Up	Up	W Down	38/38	No	SF	(1,1)
			P Up	2092/2092	Yes	SF	(1,1)
lsp-39	Up	Up	W Down	39/39	No	SF	(1,1)
			P Up	2094/2094	Yes	SF	(1,1)
lsp-40	Up	Up	W Down	40/40	No	SF	(1,1)
			P Up	2096/2096	Yes	SF	(1,1)
lsp-41	Up	Up	W Down	41/41	No	SF	(1,1)
			P Up	2098/2098	Yes	SF	(1,1)

No. of MPLS-TP LSPs: 10							
=====							

2.21.2.2 Show Router BFD session Commands

session

Syntax **session {ipv4 | ipv6} detail [lag lag-id] lag-port port-id**
session lsp-name Lsp Name
session lsp-rsvp {head | tail}
session src ip-address/link-local address dest ip-address | link-local address detail lsp-rsvp {head | tail} tunnel-id tunnel-id lsp-id lsp-id
session mpls-tp
session lsp-name Lsp Name [link-type {cc-only | cc-cv}] detail
session p2mp-interface interface-name detail
session src ip-address/link-local address detail lsp-rsvp {head | tail} rsvp-session-name [256 chars max]
session [src ip-address/link-local address] [ipv4 | ipv6]
session src ip-address/link-local address dest ip-address | link-local address
session src ip-address/link-local address detail
session summary
session type type [ipv4 | ipv6]

Context	show>router>bfd
Description	This command display BFD session information.
Parameters	ipv4 ipv6 — Displays session information for IPv4 or IPv6. detail — Displays detailed information. lag lag-id — Displays information about the specified LAG ID. lag-port port-id — Displays information about the specified LAG port ID. lsp-name Lsp Name — Displays information about the specified LSP name. lsp-rsvp {head tail} — Displays a summary of all head or tail RSVP LSP BFD sessions. The information in this show command should be the same as the show>router>bfd>session commands, but filtered for sessions associated with RSVP LSPs. The protocol field of the output should indicate lsp-rsvp. src ip-address/link-local address — Displays information about the specified source IP address or link local address. dest ip-address link-local address — Displays information about the specified destination IP address or link local address. tunnel-id tunnel-id — Displays information about the specified tunnel. lsp-id lsp-id — Displays information about the specified LSP. p2mp-interface interface-name — Displays information about the specified P2MP interface rsvp-session-name [256 chars max] — Displays a summary of all head or tail RSVP LSP BFD sessions. The information in this show command should be the same as the show>router>bfd>session commands, but filtered for sessions associated with RSVP LSPs. The protocol field of the output should indicate lsp-rsvp. type type — Specifies the type. Values iom, central, cpm-np summary — Displays summarized information.

2.21.2.3 Show RSVP Commands

interface

Syntax	interface [ip-int-name ip-address] statistics [detail]
Context	show>router>rsvp
Description	This command shows RSVP interfaces.

Parameters

ip-int-name — The name of the network IP interface. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

ip-address — The system or network interface IP address.

statistics — Displays the RSVP interface name and counts of various RSVP packets sent and received on the interface.

detail — Displays detailed information.

Output RSVP Interface Output

[Table 25](#) describes RSVP interface output fields.

Table 25 RSVP Interface Output Fields

Label	Description
Interface	The name of the IP interface.
Total Sessions	The total number of RSVP 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 sessions on this interface.
Total BW (Mbps)	The amount of bandwidth in Mb/s available to be reserved for the RSVP protocol on the interface.
Resv BW (Mbps)	The amount of bandwidth in Mb/s reserved on this interface. A value of zero (0) indicates that no bandwidth is reserved.
Adm	Down — The RSVP interface is administratively disabled. Up — The RSVP interface is administratively enabled.
Bfd	Yes — BFD is enabled on the RSVP interface. No — BFD is disabled on the RSVP interface.
Opr	Down — The RSVP interface is operationally down. Up — The RSVP interface is operationally up.
Port ID	Specifies the physical port bound to the interface.
Active Resvs	The total number of active RSVP sessions that have reserved bandwidth.
Subscription	Specifies the percentage of the link bandwidth that RSVP can use for reservation. When the value is zero (0), no new sessions are permitted on this interface.
Port Speed	Specifies the speed for the interface.
Unreserved BW	Specifies the amount of unreserved bandwidth.

Table 25 **RSVP Interface Output Fields (Continued)**

Label	Description (Continued)
Reserved BW	Specifies the amount of bandwidth in Mb/s reserved by the RSVP session on this interface. A value of zero (0) indicates that no bandwidth is reserved.
Total BW	Specifies the amount of bandwidth in Mb/s available to be reserved for the RSVP protocol on this interface.
Aggregate	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	Specifies the length of time (in s) between the hello packets that the router sends on the interface. This value must be the same for all routers attached to a common network. When the value is zero (0), the sending of hello messages is disabled.
Refresh Time	Specifies the interval between the successive Path and Resv refresh messages. RSVP declares the session down after it misses $((\text{keep-multiplier} + 0.5) * 1.5 * \text{refresh-time})$ consecutive refresh messages.
Hello Timeouts	The total number of hello messages that timed out on this RSVP interface.
Neighbors	The IP address of the RSVP neighbor.
Sent	The total number of error free RSVP packets that have been transmitted on the RSVP interface.
Recd	The total number of error free RSVP packets received on the RSVP interface.
Total Packets	The total number of RSVP packets, including errors, received on the RSVP interface.
Bad Packets	The total number of RSVP packets with errors transmitted on the RSVP interface.
Paths	The total number of RSVP PATH messages received on the RSVP interface.
Path Errors	The total number of RSVP PATH ERROR messages transmitted on the RSVP interface.
Path Tears	The total number of RSVP PATH TEAR messages received on the RSVP interface.
Resvs	The total number of RSVP RESV messages received on the RSVP interface.
Resv Confirms	The total number of RSVP RESV CONFIRM messages received on the RSVP interface.
Resv Errors	Total RSVP RESV ERROR messages received on RSVP interface.
Resv Tears	Total RSVP RESV TEAR messages received on RSVP interface.
Refresh Summaries	Total RSVP RESV summary refresh messages received on interface.
Refresh Acks	Total RSVP RESV acknowledgement messages received when refresh reduction is enabled on the RSVP interface.
Bundle Packets	Total RSVP RESV bundled packets received on the RSVP interface.

Table 25 **RSVP Interface Output Fields (Continued)**

Label	Description (Continued)
Hellos	Total RSVP RESV HELLO REQ messages received on the interface.

Sample Output

```
*A:Dut-A>config>router>mpls>lsp$ /show router rsvp interface "ip-10.10.1.1" detail
```

```
=====
RSVP Interface (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   : 1
Total Sessions    : 1
Subscription      : 100 %
Total BW          : 100 Mbps
Hello Interval    : n/a
Authentication    : Disabled
Auth Rx Seq Num   : n/a
Auth Tx Seq Num   : n/a
Refresh Reduc.    : Disabled
Bfd Enabled       : n/a
ImplicitNullLabel : Disabled*
Oper State        : Up
Active Resvs      : 0
Port Speed        : 100 Mbps
Aggregate         : Dsabl
Hello Timeouts    : n/a
Auth Key Id       : n/a
Auth Win Size     : n/a
Reliable Deli.    : Disabled
Graceful Shut.    : Disabled
GR helper         : n/a

Percent Link Bandwidth for Class Types*
Link Bw CT0       : 100
Link Bw CT1       : 0
Link Bw CT2       : 0
Link Bw CT3       : 0
Link Bw CT4       : 0
Link Bw CT5       : 0
Link Bw CT6       : 0
Link Bw CT7       : 0

Bandwidth Constraints for Class Types (Kbps)
BC0               : 100000
BC1               : 0
BC2               : 0
BC3               : 0
BC4               : 0
BC5               : 0
BC6               : 0
BC7               : 0

Bandwidth for TE Class Types (Kbps)
TE0-> Resv. Bw    : 0
TE1-> Resv. Bw    : 0
TE2-> Resv. Bw    : 0
TE3-> Resv. Bw    : 0
TE4-> Resv. Bw    : 0
TE5-> Resv. Bw    : 0
TE6-> Resv. Bw    : 0
TE7-> Resv. Bw    : 0
Unresv. Bw       : 100000
Unresv. Bw       : 100000
Unresv. Bw       : 100000
Unresv. Bw       : 100000
Unresv. Bw       : 100000
Unresv. Bw       : 100000
Unresv. Bw       : 100000

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
Neighbors        : 10.20.1.2
* indicates inherited values
=====
*A:Dut-A>config>router>mpls>lsp$
```

```
*A:SRU4>show>router>rsvp# interface
=====
RSVP Interfaces
=====
```

Interface	Total Sessions	Active Sessions	Total BW (Mbps)	Resv BW (Mbps)	Adm	Opr
system	-	-	-	-	Up	Up
aps-1	0	0	6012	0	Up	Up
aps-2	0	0	6010	0	Up	Up
aps-3	0	0	6010	0	Up	Up
sr4-1	0	0	6010	0	Up	Up
ess-7-1	9	9	100	0	Up	Up
ess-7-2	7	7	100	0	Up	Up
ess-7-3	4	4	100	0	Up	Up
ess-7-4	0	0	800	0	Up	Up
ess-7-5	0	0	800	0	Up	Up
ess-7-6	0	0	800	0	Up	Up
hubA	0	0	100	0	Up	Up
germ-1	0	0	1000	0	Up	Up
src-1.1	3	3	100	0	Up	Up
src-1.2	2	2	100	0	Up	Up
src-1.3	3	3	100	0	Up	Up
src-1.4	5	5	100	0	Up	Up
...						
g7600	0	0	1000	0	Up	Up
m160	481	481	1000	82	Up	Up

```
-----
Interfaces : 35
=====
```

```
*A:SRU4>show>router>rsvp#
=====
*A:SRU4>show>router>rsvp# interface statistics
=====
RSVP Interface (statistics)
=====
```

```
-----
Interface system
-----
```

Interface	(Sent)	(Recd.)
Total Packets	: 0	: 0
Bad Packets	: 0	: 0
Paths	: 0	: 0
Path Errors	: 0	: 0
Path Tears	: 0	: 0
Resvs	: 0	: 0
Resv Confirms	: 0	: 0
Resv Errors	: 0	: 0
Resv Tears	: 0	: 0
Refresh Summaries	: 0	: 0

```

Refresh Acks      (Sent) : 0                (Recd.): 0
Bundle Packets    (Sent) : 0                (Recd.): 0
Hellos            (Sent) : 0                (Recd.): 0
Auth Errors       (Sent) : 0                (Recd.): 0
-----
...
-----

```

Interface m160

```

-----
Interface          : Up
Total Packets      (Sent) : 883643          (Recd.): 3052503
Bad Packets        (Sent) : 0                (Recd.): 0
Paths              (Sent) : 592153          (Recd.): 373610
Path Errors        (Sent) : 464              (Recd.): 30716
Path Tears         (Sent) : 29563           (Recd.): 3480
Resvs              (Sent) : 93970           (Recd.): 2518660
Resv Confirms      (Sent) : 0                (Recd.): 0
Resv Errors        (Sent) : 136815          (Recd.): 54115
Resv Tears         (Sent) : 13338           (Recd.): 71922
Refresh Summaries  (Sent) : 0                (Recd.): 0
Refresh Acks       (Sent) : 0                (Recd.): 0
Bundle Packets     (Sent) : 0                (Recd.): 0
Hellos             (Sent) : 17340           (Recd.): 0
Auth Errors        (Sent) : 0                (Recd.): 0
=====

```

```
*A:SRU4>show>router>rsvp#
```

```
*A:SRU4>show>router>rsvp# interface "sr4-1" statistics
```

```
=====
RSVP Interface : sr4-1 (statistics)
=====
-----

```

Interface sr4-1

```

-----
Interface          : Up
Total Packets      (Sent) : 33100           (Recd.): 20405
Bad Packets        (Sent) : 0                (Recd.): 0
Paths              (Sent) : 0                (Recd.): 1833
Path Errors        (Sent) : 1783            (Recd.): 9
Path Tears         (Sent) : 0                (Recd.): 1157
Resvs              (Sent) : 76               (Recd.): 0
Resv Confirms      (Sent) : 0                (Recd.): 0
Resv Errors        (Sent) : 0                (Recd.): 0
Resv Tears         (Sent) : 1                (Recd.): 0
Refresh Summaries  (Sent) : 4                (Recd.): 33
Refresh Acks       (Sent) : 1520            (Recd.): 4
Bundle Packets     (Sent) : 0                (Recd.): 0
Hellos             (Sent) : 29716           (Recd.): 17369
Auth Errors        (Sent) : 0                (Recd.): 0
=====

```

```
*A:SRU4>show>router>rsvp#
```

```
*A:SRU4>show>router>rsvp# interface detail
```

```
=====
RSVP Interfaces (Detailed)
=====
-----

```

```
Interface : system
```

```

-----
Interface          : system
Port ID            : system
Admin State        : Up                Oper State         : Up
Active Sessions    : 0                Active Resvs        : 0
Total Sessions     : 0
Subscription       : 100 %            Port Speed          : 0 Mbps
Total BW           : 0 Mbps            Aggregate           : Dsabl
Hello Interval     : 3000 ms           Hello Timeouts      : 0
Authentication     : Disabled
Auth Rx Seq Num    : n/a              Auth Key Id         : n/a
Auth Tx Seq Num    : n/a              Auth Win Size       : n/a
Refresh Reduc.     : Enabled           Reliable Deli.      : Disabled
Bfd Enabled        : No                Graceful Shut.      : Disabled
ImplicitNullLabel  : Disabled*

```

Percent Link Bandwidth for Class Types*

```

Link Bw CT0        : 100              Link Bw CT4         : 0
Link Bw CT1        : 0                Link Bw CT5         : 0
Link Bw CT2        : 0                Link Bw CT6         : 0
Link Bw CT3        : 0                Link Bw CT7         : 0

```

Bandwidth Constraints for Class Types (Kbps)

```

BC0                : 0                BC4                 : 0
BC1                : 0                BC5                 : 0
BC2                : 0                BC6                 : 0
BC3                : 0                BC7                 : 0

```

Bandwidth for TE Class Types (Kbps)

```

TE0-> Resv. Bw     : 0                Unresv. Bw          : 0
TE1-> Resv. Bw     : 0                Unresv. Bw          : 0
TE2-> Resv. Bw     : 0                Unresv. Bw          : 0
TE3-> Resv. Bw     : 0                Unresv. Bw          : 0
TE4-> Resv. Bw     : 0                Unresv. Bw          : 0
TE5-> Resv. Bw     : 0                Unresv. Bw          : 0
TE6-> Resv. Bw     : 0                Unresv. Bw          : 0
TE7-> Resv. Bw     : 0                Unresv. Bw          : 0

```

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.

```

Interface : m160

```

-----
Interface          : m160
Port ID            : 3/2/1
Admin State        : Up                Oper State         : Up
Active Sessions    : 218              Active Resvs        : 0
Total Sessions     : 517
Subscription       : 1000 %           Port Speed          : 100 Mbps
Total BW           : 1000 Mbps         Aggregate           : Dsabl
Hello Interval     : 3000 ms           Hello Timeouts      : 0
Authentication     : Disabled
Auth Rx Seq Num    : n/a              Auth Key Id         : n/a
Auth Tx Seq Num    : n/a              Auth Win Size       : n/a
Refresh Reduc.     : Enabled           Reliable Deli.      : Disabled

```

```
Bfd Enabled      : No
ImplicitNullLabel : Disabled*
Graceful Shut.   : Disabled

Percent Link Bandwidth for Class Types*
Link Bw CT0      : 100
Link Bw CT1      : 0
Link Bw CT2      : 0
Link Bw CT3      : 0
Link Bw CT4      : 0
Link Bw CT5      : 0
Link Bw CT6      : 0
Link Bw CT7      : 0

Bandwidth Constraints for Class Types (Kbps)
BC0              : 1000000
BC1              : 0
BC2              : 0
BC3              : 0
BC4              : 0
BC5              : 0
BC6              : 0
BC7              : 0

Bandwidth for TE Class Types (Kbps)
TE0-> Resv. Bw   : 0
TE1-> Resv. Bw   : 0
TE2-> Resv. Bw   : 0
TE3-> Resv. Bw   : 0
TE4-> Resv. Bw   : 0
TE5-> Resv. Bw   : 0
TE6-> Resv. Bw   : 0
TE7-> Resv. Bw   : 0
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000
Unresv. Bw      : 1000000

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
Neighbors         : 10.100.30.20
* indicates inherited values
=====
*A:SRU4>show>router>rsvp#
```

neighbor

- Syntax** neighbor [*ip-address*] [*detail*]
- Context** show>router>rsvp
- Description** This command shows neighbor information.
- Parameters** *ip-address* — Displays RSVP information about the specified IP address.
detail — Displays detailed information.
- Output**

```
Sample Output
*A:Dut-A>config>router>mpls>lsp$ /show router rsvp neighbor
=====
RSVP Neighbors
=====
```

```

Legend :
    LR - Local Refresh Reduction          RR - Remote Refresh Reduction
    LD - Local Reliable Delivery          RM - Remote Node supports Message ID
    LG - Local Graceful Restart           RG - Remote Graceful Restart
=====
Neighbor      Interface      Hello  Last Oper  Flags
                                Change
=====
10.20.1.2      ip-10.10.1.1      N/A   0d 00:00:44
10.20.1.3      ip-10.10.2.1      N/A   0d 00:00:44
-----
Neighbors : 2
-----
*A:Dut-A>config>router>mpls>lsp$

*A:SR1# show router rsvp neighbor detail
=====
RSVP Neighbors (Detailed)
=====
Legend :
    LR - Local Refresh Reduction          RR - Remote Refresh Reduction
    LD - Local Reliable Delivery          RM - Remote Node supports Message ID
    LG - Local Graceful Restart           RG - Remote Graceful Restart
=====
Neighbor : 30.30.30.2
-----
Interface      : int_SR1_SR3      Hello State      : Up
Last Oper Change : 0d 00:01:02    Flags            :
Source Instance  : 0x6c8b7       Dst. Instance    : 0x530f8e0
Hello Refresh Time : 2 secs      Hello Timeout Time : 8 secs
Inst. Mismatch Cnt : 0          Epoch Num Rx      : 0
Srefresh Time Rem. : 0 secs      Out of order Msgs : 0
Max Msg Id Rx     : 0            GR Helper         : Disabled
Retransmitted Msgs : 0           GR Helper State    : None
GR Proc Invoked Cnt: 0           GR Nbr Restart Cap : N/A
GR Helper Time Rem : N/A         GR Nbr Recvry Time : N/A
=====

*B:edge13# show router rsvp neighbor
=====
RSVP Neighbors
=====
Legend :
    LR - Local Refresh Reduction          RR - Remote Refresh Reduction
    LD - Local Reliable Delivery          RM - Remote Node supports Message ID
    LG - Local Graceful Restart           RG - Remote Graceful Restart
=====
Neighbor      Interface      Hello  Last Oper  Flags Change
=====
10.11.101.2    e13c2_1      Up     1d 00:52:56  LR RR LD RM

LG RG
10.11.102.2    e13c2_2      Up     1d 00:52:56  LR RR LD RM
10.11.103.3    e13s1_1      Up     1d 00:52:54  LR RR LD RM

LG

```

```
10.11.104.3      e13s1_2      Up      1d 00:52:56
10.11.105.4      e13s2_1      Up      1d 00:52:56
10.11.106.4      e13s2_2      Up      1d 00:52:56
-----
Neighbors : 6
```

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 session information.
- Parameters** **session** *session-type* — Specifies the session type.

Values originate, transit, terminate, detour, detour-transit, detour-terminate, bypass-tunnel

from *ip-address* — Specifies the IP address of the originating router.

to *ip-address* — Specifies the IP address of the egress router.

lsp-name *name* — Specifies the name of the LSP used in the path.

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** RSVP Session Output

Table 26 describes RSVP session output fields.

Table 26 RSVP Session Output Fields

Label	Description
From	The IP address of the originating router.
To	The IP address of the egress router.
Tunnel ID	The IP address of the tunnel's ingress node supporting this RSVP session.
LSP ID	The ID assigned by the agent to this RSVP session.
Name	The administrative name assigned to the RSVP session by the agent.
State	Down — The operational state of this RSVP session is down. Up — The operational state of this RSVP session is up.

Sample Output

```
*A:SRU4>show>router>rsvp#    session
=====
RSVP Sessions
=====
From          To          Tunnel LSP   Name                               State
              ID      ID
-----
110.20.1.5    110.20.1.4    18     27648 b4-1::b4-1                      Up
110.20.1.5    110.20.1.4     1     37902 gsr::gsr                        Up
110.20.1.5    10.20.1.22    11     53760 to_10_20_1_22_cspf::to_10_2* Up
110.20.1.4    10.20.1.20    146    17920 to_10_20_1_20_cspf_3::to_10* Up
110.20.1.4    10.20.1.20    145    34816 to_10_20_1_20_cspf_2::to_10* Up
110.20.1.4    10.20.1.20    147    45056 to_10_20_1_20_cspf_4::to_10* Up
110.20.1.4    10.20.1.20    148    6656  to_10_20_1_20_cspf_5::to_10* Up
110.20.1.4    10.20.1.20    149    58880 to_10_20_1_20_cspf_6::to_10* Up
110.20.1.4    10.20.1.20    150    13312 to_10_20_1_20_cspf_7::to_10* Up
110.20.1.4    10.20.1.20    152    40448 to_10_20_1_20_cspf_9::to_10* Up
110.20.1.4    10.20.1.20    154    27648 to_10_20_1_20_cspf_11::to_1* Up
110.20.1.4    10.20.1.20    155    12288 to_10_20_1_20_cspf_12::to_1* Up
110.20.1.4    10.20.1.20    151    46080 to_10_20_1_20_cspf_8::to_10* Up
110.20.1.4    10.20.1.20    153    512   to_10_20_1_20_cspf_10::to_1* Up
110.20.1.4    10.20.1.22    164    62464 to_10_20_1_22_cspf_2::to_10* Up
110.20.1.4    10.20.1.20    156    37888 to_10_20_1_20_cspf_13::to_1* Up
110.20.1.4    10.20.1.20    157    24064 to_10_20_1_20_cspf_14::to_1* Up
110.20.1.4    10.20.1.20    158    19968 to_10_20_1_20_cspf_15::to_1* Up
110.20.1.4    10.20.1.20    161    59904 to_10_20_1_20_cspf_18::to_1* Up
...
110.20.1.3    110.20.1.4    54     23088 to_110_20_1_4_cspf_4::to_11* Up
-----
Sessions : 1976
=====
* indicates that the corresponding row element may have been truncated.
*A:SRU4>show>router>rsvp#

A:ALA-12# show router rsvp session lsp-name A_C_2::A_C_2 status up
=====
RSVP Sessions
=====
From          To          Tunnel LSP   Name                               State
              ID      ID
-----
10.20.1.1     10.20.1.3     2       40   A_C_2::A_C_2                      Up
-----
Sessions : 1
=====
A:ALA-12#

*A:SRU4>show>router>rsvp#    session detail
=====
RSVP Sessions (Detailed)
=====
-----
LSP : b4-1::b4-1
-----
From          : 110.20.1.5          To          : 110.20.1.4
```

```

Tunnel ID      : 18                      LSP ID        : 27648
Style          : FF                      State         : Up
Session Type   : Terminate
In Interface   : 3/2/1                  Out Interface  : n/a
In Label       : 131071                 Out Label     : n/a
Previous Hop   : 10.100.30.20           Next Hop      : n/a
SetupPriority   : 7                     Hold Priority   : 0
Class Type     : 0
SubGrpOrig ID  : 0                     SubGrpOrig Addr: 0.0.0.0
P2MP ID        : 0

```

```

Path Recd      : 7497                   Path Sent      : 0
Resv Recd      : 0                     Resv Sent      : 1757

```

Summary messages:

```

SPath Recd     : 0                     SPath Sent     : 0
SResv Recd     : 0                     SResv Sent     : 0

```

```

-----
LSP : gsr::gsr
-----

```

```

From           : 110.20.1.5             To           : 110.20.1.4
Tunnel ID      : 1                      LSP ID       : 37902
Style          : FF                      State        : Up
Session Type   : Terminate
In Interface   : 3/2/7                  Out Interface  : n/a
In Label       : 128547                 Out Label     : n/a
Previous Hop   : 160.60.60.2           Next Hop      : n/a
SetupPriority   : 7                     Hold Priority   : 0
Class Type     : 0
SubGrpOrig ID  : 0                     SubGrpOrig Addr: 0.0.0.0
P2MP ID        : 0

```

```

Path Recd      : 5225                   Path Sent      : 0
Resv Recd      : 0                     Resv Sent      : 1741

```

Summary messages:

```

SPath Recd     : 0                     SPath Sent     : 0
SResv Recd     : 0                     SResv Sent     : 0

```

```

...
-----

```

```

From           : 110.20.1.3             To           : 110.20.1.4
Tunnel ID      : 54                      LSP ID       : 23088
Style          : SE                      State        : Up
Session Type   : Terminate
In Interface   : aps-1                  Out Interface  : n/a
In Label       : 130409                 Out Label     : n/a
Previous Hop   : 104.104.0.3           Next Hop      : n/a
SetupPriority   : 7                     Hold Priority   : 0
Class Type     : 0
SubGrpOrig ID  : 0                     SubGrpOrig Addr: 0.0.0.0
P2MP ID        : 0

```

```

Path Recd      : 1                     Path Sent      : 0
Resv Recd      : 0                     Resv Sent      : 1

```

Summary messages:

```

SPath Recd     : 840                   SPath Sent     : 0
SResv Recd     : 0                     SResv Sent     : 850

```



```

=====
*A:SRU4>show>router

*A:Dut-B# show router rsvp session detour detail
=====
RSVP Sessions (Detailed)
=====
LSP : tof919::1_detour
-----
From           : 10.20.1.2           To           : 10.20.1.4
Tunnel ID      : 919                 LSP ID       : 15441
Style          : SE                  State         : Up
Session Type   : Originate (Detour)
In Interface   : n/a                 Out Interface : 1/1/2:1
In Label       : n/a                 Out Label     : 129865
Previous Hop   : n/a                 Next Hop      : 10.10.101.4
SetupPriority  : 4                   Hold Priority  : 4
Class Type     : 5
SugGrpOrig ID : 0                   SubGrpOrig Addr: 0.0.0.0
P2MP ID       : 0

Path Recd      : 0                   Path Sent     : 106
Resv Recd      : 113                Resv Sent     : 0

Summary messages:
SPath Recd     : 0                   SPath Sent    : 0
SResv Recd     : 0                   SResv Sent    : 0
=====
*A:Dut-B#

*A:Dut-C# show router mpls lsp transit detail
=====
MPLS LSPs (Transit) (Detail)
-----
LSP tof1::sec2
-----
From           : 10.20.1.2           To           : 10.20.1.4
State          : Up
SetupPriority   : 5                   Hold Priority  : 5
Class Type     : 5
In Interface   : lag-1:0             In Label      : 131068
Out Interface   : 2/1/2              Out Label     : 131068
Previous Hop   : 10.10.12.2         Next Hop      : 10.10.11.4
Reserved BW    : 1000 Kbps
=====
*A:Dut-C#

*A:Dut-B# show router rsvp session detour-terminate detail
=====
RSVP Sessions (Detailed)
=====
LSP : tof878::1_detour
-----
From           : 10.20.1.2           To           : 10.20.1.4
Tunnel ID      : 878                 LSP ID       : 14929
Style          : SE                  State         : Up

```

```
Session Type      : Terminate (Detour)
In Interface      : lag-1:0
Out Interface     : 1/1/2:8
In Label          : 131069
Out Label         : 127951
Previous Hop      : 10.10.12.3
Next Hop          : 10.10.108.4
SetupPriority     : 4
Hold Priority     : 4
Class Type       : 5
SugGrpOrig ID    : 0
SubGrpOrig Addr  : 0.0.0.0
P2MP ID          : 0
```

```
Path Recd        : 128
Path Sent         : 0
Resv Recd         : 125
Resv Sent         : 124
```

Summary messages:

```
SPath Recd       : 0
SPath Sent       : 0
SResv Recd       : 0
SResv Sent       : 0
```

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

```
*A:Dut-B# show router rsvp session bypass-tunnel detail
```

===== RSVP Sessions (Detailed)

```
=====
LSP : bypass-link10.10.108.4
```

```
-----
From          : 10.20.1.2
Tunnel ID     : 4003
Style         : FF
Session Type  : Bypass Tunnel
In Interface  : n/a
Out Interface : 1/1/2:9
In Label      : n/a
Out Label     : 124069
Previous Hop  : n/a
Next Hop      : 10.10.109.4
SetupPriority : 7
Hold Priority  : 0
Class Type   : 0
SugGrpOrig ID : 0
SubGrpOrig Addr: 0.0.0.0
P2MP ID      : 0
```

```
Path Recd        : 0
Path Sent         : 3
Resv Recd         : 4
Resv Sent         : 0
```

Summary messages:

```
SPath Recd       : 0
SPath Sent       : 0
SResv Recd       : 0
SResv Sent       : 0
=====
```

```
*A:Dut-B#
```

```
*A:Dut-B# show router rsvp session detour detail
```

===== RSVP Sessions (Detailed)

```
-----
LSP : tof919::1_detour
```

```
-----
From          : 10.20.1.2
Tunnel ID     : 919
Style         : SE
Session Type  : Originate (Detour)
In Interface  : n/a
Out Interface : 1/1/2:1
In Label      : n/a
Out Label     : 129865
Previous Hop  : n/a
Next Hop      : 10.10.101.4
```

SetupPriority	: 4	Hold Priority	: 4
Class Type	: 5		
SugGrpOrig ID	: 0	SubGrpOrig Addr:	0.0.0.0
P2MP ID	: 0		

Path Recd	: 0	Path Sent	: 106
Resv Recd	: 113	Resv Sent	: 0

Summary messages:

SPath Recd	: 0	SPath Sent	: 0
SResv Recd	: 0	SResv Sent	: 0

=====

*A:Dut-B#

*A:Dut-B# show router rsvp session detour-transit detail

=====

RSVP Sessions (Detailed)

LSP : tof919::1_detour

From	: 10.20.1.2	To	: 10.20.1.4
Tunnel ID	: 919	LSP ID	: 15441
Style	: SE	State	: Up
Session Type	: Transit (Detour)		
In Interface	: lag-1:0	Out Interface	: 1/1/2:6
In Label	: 131071	Out Label	: 127952
Previous Hop	: 10.10.12.3	Next Hop	: 10.10.106.4
SetupPriority	: 4	Hold Priority	: 4
Class Type	: 5		
SugGrpOrig ID	: 0	SubGrpOrig Addr:	0.0.0.0
P2MP ID	: 0		

Path Recd	: 119	Path Sent	: 123
Resv Recd	: 121	Resv Sent	: 120

Summary messages:

SPath Recd	: 0	SPath Sent	: 0
SResv Recd	: 0	SResv Sent	: 0

=====

*A:Dut-B#

*A:Dut-B# show router rsvp session detour-terminate detail

=====

RSVP Sessions (Detailed)

LSP : tof878::1_detour

From	: 10.20.1.2	To	: 10.20.1.4
Tunnel ID	: 878	LSP ID	: 14929
Style	: SE	State	: Up
Session Type	: Terminate (Detour)		
In Interface	: lag-1:0	Out Interface	: 1/1/2:8
In Label	: 131069	Out Label	: 127951
Previous Hop	: 10.10.12.3	Next Hop	: 10.10.108.4
SetupPriority	: 4	Hold Priority	: 4
Class Type	: 5		
SugGrpOrig ID	: 0	SubGrpOrig Addr:	0.0.0.0

```

P2MP ID          : 0

Path Recd        : 128          Path Sent       : 0
Resv Recd        : 125          Resv Sent       : 124

Summary messages:
SPath Recd       : 0            SPath Sent      : 0
SResv Recd       : 0            SResv Sent      : 0
=====
*A:Dut-B#

*A:Dut-B# show router rsvp session bypass-tunnel detail
=====
RSVP Sessions (Detailed)
-----
LSP : bypass-link10.10.108.4
-----
From          : 10.20.1.2          To          : 10.10.109.4
Tunnel ID     : 4003              LSP ID      : 6
Style         : FF                State        : Up
Session Type  : Bypass Tunnel
In Interface  : n/a              Out Interface : 1/1/2:9
In Label      : n/a              Out Label    : 124069
Previous Hop  : n/a              Next Hop     : 10.10.109.4
SetupPriority : 7                 Hold Priority : 0
Class Type    : 0
SugGrpOrig ID : 0                 SubGrpOrig Addr: 0.0.0.0
P2MP ID      : 0

Path Recd     : 0          Path Sent      : 3
Resv Recd     : 4          Resv Sent      : 0

Summary messages:
SPath Recd    : 0          SPath Sent     : 0
SResv Recd    : 0          SResv Sent     : 0
=====
*A:Dut-B#

```

statistics

Syntax	statistics
Context	show>router>rsvp
Description	This command displays global statistics in the RSVP instance.
Output	RSVP Statistics Output

[Table 27](#) describes RSVP statistics output fields.

Table 27 **RSVP Statistics Output Fields**

Label	Description
PATH Timeouts	The total number of path timeouts.
RESV Timeouts	The total number of RESV timeouts.

Sample Output

```
*A:SR1# /show router rsvp statistics
=====
RSVP Global Statistics
=====
PATH Timeouts      : 0                RESV Timeouts      : 0
GR Helper PATH Tim*: 0                GR Helper RESV Tim*: 0
=====
* indicates that the corresponding row element may have been truncated.

*A:SRU4>show>router>rsvp# statistics
=====
RSVP Global Statistics
=====
PATH Timeouts      : 1026             RESV Timeouts      : 182
=====
*A:SRU4>show>router>rsvp#
```

status

Syntax	status
Context	show>router>rsvp
Description	This command displays RSVP status.
Output	RSVP Status

[Table 28](#) describes RSVP status output fields.

Table 28 **RSVP Status Output Fields**

Label	Description
Admin Status	Down — RSVP is administratively disabled. Up — RSVP is administratively enabled.
Oper Status	Down — RSVP is operationally down. Up — RSVP is operationally up.

Table 28 **RSVP Status Output Fields (Continued)**

Label	Description (Continued)
Keep Multiplier	Displays the keep-multiplier <i>number</i> used by RSVP to declare that a reservation is down or the neighbor is down.
Refresh Time	Displays the refresh-time interval (in s), between the successive Path and Resv refresh messages.
Message Pacing	Enabled — RSVP messages, specified in the max-burst command, are sent in a configured interval, specified in the period command. Disabled — Message pacing is disabled. RSVP message transmission is not regulated.
Pacing Period	Displays the time interval (in ms), when the router can send the specified number of RSVP messages specified in the rsvp max-burst command.
Max Packet Burst	Displays the maximum number of RSVP messages that are sent in the specified period under normal operating conditions.
Soft Preemption Timer	Displays the time (in s), a node holds on to a reservation for which it has triggered the soft preemption procedure.
Rapid Retransmit	Displays the value of the rapid retransmission interval.
Rapid Retry Limit	Displays the rapid retry limit.
Graceful Shutdown	Specifies whether graceful shutdown of the RSVP node is enabled.

Sample Output

```

B:# 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
Implicit Null Label: Disabled      Node-id in RRO   : Exclude
P2P Merge Point Ab*: 10          P2MP Merge Point A*: 10
DiffServTE AdmModel: Basic
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
=====
* indicates that the corresponding row element may have been truncated.

=====
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 RRO   : Exclude
P2P Merge Point Ab*: Disabled     P2MP Merge Point A*: Disabled
DiffServTE AdmModel: Basic        Entropy Label    : Enabled
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
=====
* indicates that the corresponding row element may have been truncated.

```

lsp-bfd

Syntax **lsp-bfd**
lsp-bfd local-bfd-discrim *bfd-discriminator*
lsp-bfd lsp-name *lsp-name*

Context show>test-oam

Description This command displays information about Bidirectional Forwarding Detection (BFD) sessions on LSPs.

Parameters **local-bfd-discrim** *bfd-discriminator* — Displays information about the unique local discriminator for this session.

lsp-name *lsp-name* — Specifies information about the specified LSP name.

Output

Sample Output

```
A:bkvm1# show test-oam lsp-bfd local-bfd-discrim 1
-----
LSP Ping Bootstrap and Periodic Verification Information for BFD on an LSP
-----
LSP Name : Z
LSP Path Status      : unknown
Replying Node        : (None)
Latest Return Code    : NoRtnCode (0)
Latest Return Subcode : 3
Local BFD Discriminator : 1          Remote BFD Discriminator : 4111222333
Tx LSP Ping Requests  : 2123456789  Rx LSP Ping Requests   : 3123456789
-----
A:bkvm1#
A:bkvm1# show test-oam lsp-bfd local-bfd-discrim 1
-----
LSP Ping Bootstrap and Periodic Verification Information for BFD on an LSP
-----
LSP Name : Z
LSP Path Status      : active
Replying Node        : 240.241.242.243
Latest Return Code    : DSRtrMatchLabel (8)
Latest Return Subcode : 4
Local BFD Discriminator : 1          Remote BFD Discriminator : 4111222333
Tx LSP Ping Requests  : 2123456789  Rx LSP Ping Requests   : 3123456789
-----
A:bkvm1#
A:bkvm1# show test-oam lsp-bfd local-bfd-discrim 1
-----
LSP Ping Bootstrap and Periodic Verification Information for BFD on an LSP
-----
LSP Name : Z
LSP Path Status      : inactive
Replying Node        : f0f1:f2f3:f4f5:f6f7:f8f9:fafb:fcfd:feff
Latest Return Code    : DSRtrUnmatchLabel (10)
Latest Return Subcode : 5
Local BFD Discriminator : 1          Remote BFD Discriminator : 4111222333
Tx LSP Ping Requests  : 2123456789  Rx LSP Ping Requests   : 3123456789
-----
A:bkvm1#
A:bkvm1# show test-oam lsp-bfd local-bfd-discrim 1
-----
```



```
LSP Ping Bootstrap and Periodic Verification Information for BFD on an LSP
-----
LSP Name : Z
LSP Path Status      : unknown
Replying Node       : (None)
Latest Return Code   : DSNoMac (16)
Latest Return Subcode : 6
Local BFD Discriminator : 1           Remote BFD Discriminator : 4111222333
Tx LSP Ping Requests  : 2123456789   Rx LSP Ping Requests    : 3123456789
-----
A:bkvm1#
```

2.21.2.4 Tools Commands

bypass-tunnel

- Syntax** **bypass-tunnel** [*lsp-name*] **plr**
- Context** tools>dump>router>mpls
- Description** This command displays information about the MPLS bypass tunnel.

lspinfo

- Syntax** **lspinfo**
- Context** tools>dump>router>mpls
- Description** This command dumps LSP information for MPLS.

ftn

- Syntax** **ftn**
- Context** tools>dump>router>mpls
- Description** This command dumps FTN information for MPLS.

ilm

- Syntax** **ilm**
- Context** tools>dump>router>mpls
- Description** This command dumps ILM information for MPLS.

memory-usage

Syntax	memory-usage
Context	tools>dump>router>mpls
Description	This command dumps memory usage information for MPLS.

adjust-autobandwidth

Syntax	adjust-autobandwidth [lsp <i>lsp-name</i> [force [bandwidth <i>mbps</i>]]]
Context	tools>perform>router>mpls
Description	This command initiates an immediate automatic bandwidth adjustment attempt for either one specific LSP or all active LSPs. The automatic bandwidth adjustment is made to the primary or secondary path of the LSP, whichever is the currently active path. If an LSP is not specified, then the system assumes the command applies to all LSPs. The optional force parameter, which is available only when an LSP is referenced, determines whether adjust-up and adjust-down threshold checks are applied. If force is not specified then the maximum average data rate must differ from the current reservation by more than the adjust-up or adjust-down thresholds, otherwise no bandwidth adjustment occurs. If the force option is specified then, bandwidth adjustment ignores the configured thresholds. If a bandwidth is specified as part of the force option then the bandwidth of the LSP is changed to this specific value, otherwise the bandwidth is changed to the maximum average data rate that has been measured by the system in the current adjust interval. The adjust-count and maximum average data rate are not reset by the manual auto-bandwidth command, whether or not the bandwidth adjustment succeeds or fails. The overflow count is reset only if the manual auto-bandwidth attempt is successful.
Default	none
Parameters	<i>lsp-name</i> — The name of the LSP to which this command applies. If this parameter is not supplied the command applies to all active LSPs. Values String (32 chars max) Default none <i>mbps</i> — The bandwidth that the LSP should be immediately resized to. Values 0 to 1000000 Default none

force-switch-path

Syntax	force-switch-path [lsp <i>lsp-name</i>] [path <i>path-name</i>]

Context	tools>perform>router>mpls
Description	Use this command to move from a standby path to any other standby path regardless of priority. The no form of the command reverts to priority path.
Parameters	<i>lsp-name</i> — Specifies an existing LSP name to move. <i>path-name</i> — Specifies the path name to which to move the specified LSP.

plr

Syntax	plr
Context	tools>dump>router>mpls>bypass-tunnel
Description	Dump the Point of Local Repair (LPR) information for the MPLS bypass tunnel.
Output	

Sample Output

```
tools dump router mpls bypass-tunnel plr
=====
MPLS Bypass Tunnels
=====
Legend :  m - Manual      d - Dynamic      p - P2mp
=====
To          State  Out I/F      Out Label    Reserved    Protected    Type
              BW (Kbps)  LSP Count
-----
10.10.12.1   Up    1/1/4        124181        0           369          d

To          : 10.10.12.1      State          : Up
Out I/F     : 1/1/4          Out Label      : 124181
Up Time     : 0d 19:24:13    Active Time    : n/a
Reserved BW : 0 Kbps         Protected LSP Count : 369
Type        : Dynamic
SetupPriority : 7             Hold Priority   : 0
Class Type  : 0             Tunnel Id      : 63697
Actual Hops  :
    10.10.12.2(S)          -> 10.10.12.1(S)

Plr List: (Last PlrIdx 2)
-----
    PLR List Index = 1
    PLR current State = PLRS_CONNECTED
    NextNodeSysId = 8.8.8.8
    AvoidNodeId    = 2.2.2.2
    NodeProtect    = 2 (Node Protect)
    LSP Count      = 197
    PLR List Index = 2
    PLR current State = PLRS_BackupInUse
    NextNodeSysId = 8.8.8.8
```

```

AvoidNodeId    = 2.2.2.2
NodeProtect    = 2 (Node Protect)
LSP Count      = 203

```

cspf

Syntax	cspf to <i>ip-addr</i> [from <i>ip-addr</i>] [bandwidth <i>bandwidth</i>] [include-bitmap <i>bitmap</i>] [exclude-bitmap <i>bitmap</i>] [hop-limit <i>limit</i>] [exclude-address <i>excl-addr</i> [<i>excl-addr</i> ...(up to 8 max)]] [use-te-metric] [strict-srlg] [srlg-group <i>grp-id</i> ...(up to 8 max)] [exclude-node <i>excl-node-id</i> [<i>excl-node-id</i> ...(up to 8 max)]] [skip-interface <i>interface-name</i>] [ds-class-type <i>class-type</i>] [cspf-reqtype <i>req-type</i>] [least-fill-min-thd <i>thd</i>] [setup-priority <i>val</i>] [hold-priority <i>val</i>]
Context	tools>perform>router>mpls
Description	This command computes a CSPF path with specified user constraints.
Default	none
Parameters	to <i>ip-addr</i> — Specify the destination IP address. from <i>ip-addr</i> — Specify the originating IP address. bandwidth <i>bandwidth</i> — Specifies the amount of bandwidth in Mb/s to be reserved. include-bitmap <i>bitmap</i> — Specifies to include a bit-map that specifies a list of admin groups that should be included during setup. exclude-bitmap <i>bitmap</i> — Specifies to exclude a bit-map that specifies a list of admin groups that should be included during setup. hop-limit <i>limit</i> — Specifies the total number of hops a detour LSP can take before merging back onto the main LSP path. exclude-address <i>ip-addr</i> — Specifies IP addresses, up to 8, that should be included during setup. use-te-metric — Specifies the use of the traffic engineering metric used on the interface. strict-srlg — Specifies whether to associate the LSP with a bypass or signal a detour if a bypass or detour satisfies all other constraints except the SRLG constraints. srlg-group <i>grp-id</i> — Specifies up to 8 Shared Risk Loss Groups (SRLGs). An SRLG group represents a set of interfaces which could be subject to the same failures or defects and thus share the same risk of failing. Values 0 to 4294967295 exclude-node <i>excl-node-id</i> — specifies a list of address that should be excluded when this LSP is setup. skip-interface <i>interface-name</i> — Specifies an interface name that should be skipped during setup. ds-class-type <i>class-type</i> — Specifies the class type (CT) associated with this LSP. Values 0 to 7

cspf-reqtype *req-type* — Specifies the req. type.

Values all, random, least-fill

least-fill-min-thd *thd* — Specifies whether the use of the least-fill path selection method for the computation of the path of this LSP is enabled.

Values 1 to 100

setup-priority *val* — Specifies the setup priority to use when insufficient bandwidth is available to setup an LSP.

Values 0 to 7

hold-priority *val* — Specifies the hold priority value to use when insufficient bandwidth is available to setup an LSP.

Values 0 to 7

resignal

Syntax **resignal** {**lsp** *lsp-name* **path** *path-name* | **delay** *minutes*}
resignal {**p2mp-lsp** *p2mp-lsp-name* **p2mp-instance** *p2mp-instance-name* | **p2mp-delay** *p2mp-minutes*}
resignal {**p2mp-lsp** *p2mp-lsp-name* **p2mp-instance** *p2mp-instance-name* | **p2mp-delay** *p2mp-minutes*}

Context tools>perform>router>mpls

Description This command resignals a specific LSP path. The *minutes* parameter configures the global timer or all LSPs for resignal. If only *lsp-name* and *path-name* are provided, the LSP will be resignaled immediately.

Parameters *lsp-name* — Specifies an existing LSP name to resignal.
path-name — Specifies an existing path name to resignal.
delay *minutes* — Configures the global timer or all LSPs to resignal.
p2mp-lsp *p2mp-lsp-name* — Specifies an existing point-to-multipoint LSP name (for the
. . .
p2mp-instance *p2mp-instance-name* — Specifies a name that identifies the P2MP LSP instance
p2mp-delay *p2mp-minutes* — Specifies the delay time, in minutes.
Values 0 to 60

resignal-bypass

Syntax **resignal-bypass** {**lsp** *bypass-lsp-name* [**force**] | **delay** *minutes*}

Context tools>perform>router>mpls

Description	This command performs a manual re-optimization of a specific dynamic or manual bypass LSP, or of all dynamic bypass LSPs. The name of a manual bypass LSP is the one provided by the user at configuration time. The name of a dynamic bypass LSP is shown in the output of “show>router>mpls>bypass-tunnel dynamic detail”. The delay option triggers the global re-optimization of all dynamic bypass LSPs at the expiry of the specified delay. In essence, this option forces the global bypass resignal timer to expire after an amount of time equal to the value of the delay parameter. This option has no affect on a manual bypass LSP. However, when a specific bypass LSP name is specified, the named dynamic or manual bypass LSP is not signaled and the associations are not evaluated even if the new bypass LSP path has the same cost as the current one. This is a different behavior from that of the similar command for the primary or secondary path of an LSP as a bypass LSP can have a large number of PSB associations. In the specific case where the name corresponds to a manual bypass LSP, the LSP is torn down and re-signaled using the new path provided by CSPF if no PSB associations exist. If there is one or more PSB association but no PLR is active, the command is failed and the user is asked to explicitly enter the force option. In this case, the manual bypass LSP is torn down and re-signaled, leaving temporarily the associated LSP primary paths unprotected. Finally, if one or more PLRs associated with the manual bypass LSP is active, the command is failed.
Parameters	<i>lsp bypass-lsp-name</i> [<i>force</i>] — Specifies the name of the dynamic or manual bypass LSP. The force option is required when the name corresponds to a manual bypass LSP and the LSP has PSB associations. delay <i>minutes</i> — Specifies the time, in minutes, MPLS waits before attempting to re-signal dynamic bypass LSP paths originated on the system. Values 0 to 30

revert

Syntax	revert [<i>lsp lsp-name</i>]
Context	tools>perform>router>mpls
Description	Use this command to cause a named LSP, which is currently using a secondary path and for which the revert-timer has been configured, to switch back to using the primary path. Any outstanding revert-timer is canceled. The primary path must be up for this command to be successful.
Parameters	<i>lsp-name</i> — Specifies an existing LSP name to move.

sr-te-cspf

Syntax	sr-te-cspf <i>to ip-address</i> [path <i>path-name</i>] [from <i>ip-address</i>]
Context	tools>perform>router>mpls
Description	This command computes and returns the Segment Routing label stack for any user-specified MPLS path to a given destination router. MPLS passes the specified path information to TE-DB, which converts the list of hops into a label stack by scanning the TE database for adjacency and node SID information which belongs to the router or link identified by each hop address. If the conversion is successful, the TE database will return the actual selected hop SIDs plus labels, as well the configured path hop addresses which were used as the input for this conversion.
Parameters	to <i>ip-address</i> — Specifies the destination address. path <i>path-name</i> — Specifies the path name. from <i>ip-address</i> — Specifies the originating IP address.

switch-path

Syntax	switch-path [lsp <i>lsp-name</i>] [path <i>path-name</i>]
Context	tools>perform>router>mpls
Description	Use this command to move from a standby (or an active secondary) to another standby of the same priority. If a new standby path with a higher priority or a primary path comes up after the tools perform command is executed, the path re-evaluation command runs and the path is moved to the path specified by the outcome of the re-evaluation.
Parameters	<i>lsp-name</i> — Specifies an existing LSP name to move. <i>path-name</i> — Specifies the path name to which to move the specified LSP.

te-lspinfo

Syntax	te-lspinfo [endpoint <i>ip-address</i>] [sender <i>ip-address</i>] [lspid <i>lsp-id</i>] [detail] [p2p p2p-tid <i>tunnel-id</i>] te-lspinfo [endpoint <i>ip-address</i>] [sender <i>ip-address</i>] [lspid <i>lsp-id</i>] [detail] [p2mp p2mp-tid <i>tunnel-id</i>]{ [phops] [nhops] [s2l <i>ip-address</i>] }
Context	tools>dump>router>mpls
Description	This command displays TE LSP information for MPLS.
Default	none
Output	

Sample Output

```

B:Dut-R# tools dump router mpls te-lspinfo
Key P2P: Session(10.10.3.2, 201, 3.3.3.3) Sender(3.3.3.3, 2) PHOP(10.10.3.1), Flags
0x0

Key P2P: Session(10.10.3.1, 1035, 4.4.4.4) Sender(4.4.4.4, 22) PHOP(10.10.11.2), Fla
gs
0x0

Key P2MP: Session(0.0.0.0, 1, 4.4.4.4) Sender(4.4.4.4, 52226) PHOP(0.0.0.0) Flags 0x
10
  S2L [1] Key: endPoint to 2.2.2.2 subGroupId - 1 subGroupOrigId - 4.4.4.4
  S2L [2] Key: endPoint to 10.10.2.2 subGroupId - 3 subGroupOrigId - 4.4.4.4
  S2L [3] Key: endPoint to 10.10.13.2 subGroupId - 4 subGroupOrigId - 4.4.4.4

Key P2MP: Session(0.0.0.0, 2, 4.4.4.4) Sender(4.4.4.4, 51714) PHOP(0.0.0.0) Flags 0x
10
  S2L [1] Key: endPoint to 2.2.2.2 subGroupId - 1 subGroupOrigId - 4.4.4.4
  S2L [2] Key: endPoint to 10.10.2.2 subGroupId - 3 subGroupOrigId - 4.4.4.4
  S2L [3] Key: endPoint to 10.10.13.2 subGroupId - 4 subGroupOrigId - 4.4.4.4

Key P2MP: Session(0.0.0.0, 3, 4.4.4.4) Sender(4.4.4.4, 53250) PHOP(0.0.0.0) Flags 0x
10

*A:Dut-T# tools dump router mpls te-lspinfo p2mp-tid 102 nhops

Key P2MP: Session(0.0.0.0, 102, 4.4.4.4) Sender(4.4.4.4, 3074) PHOP(0.0.0.0) Flags
0x10
-----
List of NEXT HOPS
-----

NextHop [1] =>
Key: Nhop - isFrr 0, outIf 0, NextHop 0.0.0.0 label -
128843 global Instance 0 is
Leaf node
-----
Primary NHLFE => outLabel - 0 and NextHop - 0.0.0.0, outIf 0 (0)
Port(NONE) NhIdx 0, ProtNhIdx 0, NumS2L 1
ProtectInstance - 0, ProtectGroup 0
POP
No Backup NHLFEs for this Ltn entry
Mid List : 3428 numS2Ls - 1 (Primary MID),

NextHop [2] =>
Key: Nhop - isFrr 0, outIf 3, NextHop 10.10.13.2 label - 128806 global Instance -
48747
-----
Primary NHLFE => outLabel - 128806 and NextHop - 10.10.13.2, outIf 3 (126)
Port(9/1/1) NhIdx 4322, ProtNhIdx 2275, NumS2L 1
ProtectInstance - 1, ProtectGroup 126
SWAP
Backup NHLFE => outLabel - 130223 and NextHop - 10.10.3.2, outIf 5 (124)
Port(9/2/3) outPushLabel 128806, NhIdx 5469, ProtNhIdx 0, NumS2L 1
Mid List : 3428 numS2Ls - 1 (Primary MID),

NextHop [3] =>

```



```
Key: Nhop - isFrr 0, outIf 4, NextHop 10.10.2.2 label - 128836 global Instance -
48974
-----
Primary NHLFE => outLabel - 128836 and NextHop - 10.10.2.2, outIf 4 (125)
      Port(lag-1) NhIdx 4292, ProtNhIdx 2245, NumS2L 2
      ProtectInstance - 1, ProtectGroup 125
SWAP
Backup NHLFE => outLabel - 130223 and NextHop - 10.10.3.2, outIf 5 (124)
      Port(9/2/3) outPushLabel 128836, NhIdx 5659, ProtNhIdx 0, NumS2L 2
Mid List :      3428 numS2Ls - 1 (Primary MID),      3471 numS2Ls - 1 (Backup MID),

S2L [1] Key: endPoint to 2.2.2.2 subGroupId - 1 subGroupOrigId - 4.4.4.4
S2L [2] Key: endPoint to 3.3.3.3 subGroupId - 2 subGroupOrigId - 4.4.4.4
S2L [3] Key: endPoint to 10.10.2.2 subGroupId - 3 subGroupOrigId - 4.4.4.4
S2L [4] Key: endPoint to 10.10.13.2 subGroupId - 4 subGroupOrigId - 4.4.4.4

Total TeLspInfo Count      : 1
```

tp-interface

Syntax	tp-interface <i>interface-num</i>
Context	tools>dump>router>mpls
Description	This command displays MPLS-TP interface information.
Parameters	<i>interface-num</i> — specifies the MPLS-TP interface
	Values 1 to 4294967295
Output	

Sample Output

```
1-2 (1/1/3) : State Up, Cnt/Tm - Dn: 1/000 00:30:11.150, Up 1/000 00:30:15.160
MEP: ifNum 1 (Validate: Y), ais: Y (txCnt 0)
ais - Flags: 0x1 (R), Timers:
```

psb

Syntax	psb
Context	tools>dump>router>rsvp
Description	This command displays RSVP information.
Output	

Sample Output

```
*A:Dut-A>config>router>mpls>lsp$ /tools dump router rsvp psb detail
```

```

-----
PSB:
P2P: Session (To: 10.20.1.4 - 61441 - 10.20.1.1), Sender (10.20.1.1 - 2) PHop
255.255.255.255

PSB CurrState: BACKUPS_CONNECTED PrevState: BACKUPS_INIT Flags: 0x0
LocalLabel 0 OutLabel 131070
Incoming IfIndex: Interface: Local API(-1)
Refresh interval 0, Send Path refresh in 3 secs, Path Refresh timeout 0 secs
PrevHop: Ctype 1 Addr 255.255.255.255, LIH 0
DnStream Nbr: Addr-> 10.20.1.3 IfIndex ip-10.10.2.1(3)
UpStream Neighbor is NULLP
Session Attribute:
    Session Name: bypass-node10.20.1.2
    HoldPri: 0 SetupPri: 7 Flags: 0x2
    Ctype: 7, IncludeGroup: 0x0 IncludeAllGroup: 0x0 ExcludeGroup: 0x0
ClassType: Absent
TSpec: Flags 0x8000 QOSC 0, PDR (infinity), PBS 0.000 bps, CDR (0.000 bps) MTU: 0
CSPF Hop List: ->
    (1) UnnumIfId 3 RtrId 10.20.1.1 EgrAdmGrp 0x0 (Strict)
    (2) UnnumIfId 2 RtrId 10.20.1.3 EgrAdmGrp 0x0 (Strict)
    (3) UnnumIfId 5 RtrId 10.20.1.4 EgrAdmGrp 0x0 (Strict)
PSB RRO : ->
    (1) * Flags : 0x0 : U
    (1) * UnInf : 10.20.1.1, 3
PSB SENT RRO : ->
    (1) * Flags : 0x0 : U
    (1) * UnInf : 10.20.1.1, 3
PSB FILTERSPEC RRO : ->
    (1) * Flags : 0x0 : U
    (1) * UnInf : 10.20.1.3, 2
    (2) * Flags : 0x1 : Global
    (2) * Label : 131070
    (3) * Flags : 0x0 : U
    (3) * UnInf : 10.20.1.4, 5
    (4) * Flags : 0x1 : Global
    (4) * Label : 131070
PSB ERO : ->
    (1) Unnumbered RouterId 10.20.1.1, LinkId 3, Strict
    (2) Unnumbered RouterId 10.20.1.3, LinkId 2, Strict
    (3) Unnumbered RouterId 10.20.1.4, LinkId 5, Strict
PSB SENT ERO : ->
    (1) Unnumbered RouterId 10.20.1.3, LinkId 2, Strict
    (2) Unnumbered RouterId 10.20.1.4, LinkId 5, Strict
SendTempl: Sender:10.20.1.1_2
AdSpec Present - Flags: 0x0
AdSpec General
- Service Break bit : 0x0
- IS Hop Count : 0x0
- Path Bandwidth Estimate : 0x0
- Minimum Path latency : 0x0
- Composed path MTU : 0

Num Paths Received :0
Num Paths Transmitted:5
Num Resvs Received :8
Num Resvs Transmitted:0

Num Summary Paths Received :0

```

```

Num Summmary Paths Transmitted:0
Num Summmary Resvs Received :0
Num Summmary Resvs Transmitted:0
Created at 91359 (26 secs back)
-----
PSB:
P2P: Session (To: 10.20.1.6 - 1 - 10.20.1.1), Sender (10.20.1.1 -
30208) PHop 0.0.0.0

PSB CurrState: PRIMARYS_CONNECTED PrevState: PRIMARYS_INIT Flags: 0x8
LocalLabel 0 OutLabel 131071
Incoming IfIndex: Interface: Local API(-1)
Refresh interval 5, Send Path refresh in 4 secs, Path Refresh timeout 0 secs
PrevHop: Ctype 1 Addr 0.0.0.0, LIH 0
DnStream Nbr: Addr-> 10.20.1.2 IfIndex ip-10.10.1.1(2)
UpStream Neighbor is NULLP
Session Attribute:
    Session Name: 1::1
    HoldPri: 0 SetupPri: 7 Flags: 0x17
    Ctype: 7, IncludeGroup: 0x0 IncludeAllGroup: 0x0 ExcludeGroup: 0x0
ClassType: Absent
TSpec: Flags 0x8000 QOSC 1, PDR (infinity), PBS 0.000 bps, CDR (0.000 bps) MTU: 0
CSPF Hop List: ->
    (1) UnnumIfId 2 RtrId 10.20.1.1 EgrAdmGrp 0x0 (Strict)
    (2) UnnumIfId 2 RtrId 10.20.1.2 EgrAdmGrp 0x0 (Strict)
    (3) UnnumIfId 2 RtrId 10.20.1.4 EgrAdmGrp 0x0 (Strict)
    (4) UnnumIfId 2 RtrId 10.20.1.6 EgrAdmGrp 0x0 (Strict)
PSB RRO : ->
    (1) * Flags : 0x9 : U LP_AVAIL NODE
    (1) * UnInf : 10.20.1.1, 2
PSB SENT RRO : ->
    (1) * Flags : 0x0 : U
    (1) * UnInf : 10.20.1.1, 2
PSB FILTERSPEC RRO : ->
    (1) * Flags : 0x9 : U LP_AVAIL NODE
    (1) * UnInf : 10.20.1.2, 2
    (2) * Flags : 0x1 : Global
    (2) * Label : 131071
    (3) * Flags : 0x1 : U LP_AVAIL
    (3) * UnInf : 10.20.1.4, 2
    (4) * Flags : 0x1 : Global
    (4) * Label : 131071
    (5) * Flags : 0x0 : U
    (5) * UnInf : 10.20.1.6, 2
    (6) * Flags : 0x1 : Global
    (6) * Label : 131071
PSB ERO : ->
    (1) Unnumbered RouterId 10.20.1.2, LinkId 2, Strict
    (2) Unnumbered RouterId 10.20.1.4, LinkId 2, Strict
    (3) Unnumbered RouterId 10.20.1.6, LinkId 2, Strict
PSB SENT ERO : ->
    (1) Unnumbered RouterId 10.20.1.2, LinkId 2, Strict
    (2) Unnumbered RouterId 10.20.1.4, LinkId 2, Strict
    (3) Unnumbered RouterId 10.20.1.6, LinkId 2, Strict
SendTempl: Sender:10.20.1.1_30208
AdSpec not present
FRR: Flags 0x2 HopLimit 16 SetupPri 7 HoldPri 0 IncludeAny 0x0 ExcludeAny 0x0
IncludeAll 0x0

```

```
PLR: Flag (0x166) State PLRS_BYPASS_UP AvoidNodeId 10.20.1.2 inIntf -1 inLabel 0
PLR: FRRRequestCount: 1 CSPFFailures: 0 ProtectionType: NodeProtect

Num Paths Received    :0
Num Paths Transmitted:5
Num Resvs Received    :5
Num Resvs Transmitted:0

Num Summmary Paths Received    :0
Num Summmary Paths Transmitted:0
Num Summmary Resvs Received    :0
Num Summmary Resvs Transmitted:0
Created at 91359 (28 secs back)
-----
Total PSB Count      : 2
```

rsb

Syntax	rsb
Context	tools>dump>router>rsvp
Description	This command displays RSVP RSB information.
Output	

Sample Output

```
4)*A:Dut-A>config>router>mpls>lsp$ /tools dump router rsvp rsb detail
-----
RSB:
  EndPt 10.20.1.4  Tid 61441  XTid 10.20.1.1  Sndr 10.20.1.1  LspId 2  ifIndex 3 NHop

20.20.1.3
Style FF, refresh in 0 secs
RSVP NextHop 20.20.1.3, LIH 3 (TLV: RtrId 10.20.1.3 IntfId 2)
CT Shared Reservation Info:
No Reservation:
FlowSpec :Flags 0x8000 QOSC 1, PDR (infinity), PBS 0.000 bps, CDR (0.000 bps)
          CBS 0, EBS 0, RSpecR 0, RSpecS 0 MTU 1500 MPU 20
FwdFlowspec :Flags 0x0 QOSC 0, PDR (0.000 bps), PBS 0.000 bps, CDR (0.000 bps)
            CBS 0, EBS 0, RSpecR 0, RSpecS 0 MPU 0
FilterSpec:
Timeout in : 26 secs, LocLabel: 0  Sender: 10.20.1.1 lspId: 2 OutIfId: 0
RRO :
  (1) * Flags : 0x0 :      U
  (1) * UnInf : 10.20.1.3, 2
  (2) * Flags : 0x1 :      Global
  (2) * Label : 131070
  (3) * Flags : 0x0 :      U
  (3) * UnInf : 10.20.1.4, 5
  (4) * Flags : 0x1 :      Global
  (4) * Label : 131070
-----
-----
```

```

RSB:
  EndPt 10.20.1.6  Tid 1  XTid 10.20.1.1  Sndr 0.0.0.0  LspId 0  ifIndex 2 NHop 20.20
.1.2
Style SE, refresh in 0 secs
RSVP NextHop 20.20.1.2, LIH 2 (TLV: RtrId 10.20.1.2 IntfId 2)
CT Shared Reservation Info:
No Reservation:
FlowSpec :Flags 0x8000 QOSC 1, PDR (infinity), PBS 0.000 bps, CDR (0.000 bps)
          CBS 0, EBS 0, RSpecR 0, RSpecS 0 MTU 1496 MPU 20
FwdFlowspec :Flags 0x0 QOSC 0, PDR (0.000 bps), PBS 0.000 bps, CDR (0.000 bps)
            CBS 0, EBS 0, RSpecR 0, RSpecS 0 MPU 0
FilterSpec:
Timeout in : 21 secs, LocLabel: 0  Sender: 10.20.1.1 lspId: 30208 OutIfId: 0
RRO :
  (1) * Flags : 0x9 :      U LP_AVAIL NODE
  (1) * UnInf : 10.20.1.2, 2
  (2) * Flags : 0x1 :      Global
  (2) * Label : 131071
  (3) * Flags : 0x1 :      U LP_AVAIL
  (3) * UnInf : 10.20.1.4, 2
  (4) * Flags : 0x1 :      Global
  (4) * Label : 131071
  (5) * Flags : 0x0 :      U
  (5) * UnInf : 10.20.1.6, 2
  (6) * Flags : 0x1 :      Global
  (6) * Label : 131071
-----
Total RSB Count   : 2

```

trap-suppress

Syntax	trap-suppress <i>number-of-traps</i> <i>time-interval</i>
Context	tools>perform>router>mpls
Description	This command modifies thresholds for trap suppression. The <i>time-interval</i> parameter is used to suppress traps after a certain number of traps have been raised within the <i>time-interval</i> period of time. By executing this command, there will be no more than <i>number-of-traps</i> within <i>time-interval</i> .
Parameters	<i>number-of-traps</i> — Specifies to the number of traps raised within a period of time before suppression occurs. Values 100 to 1000, in multiples of 100 <i>time-interval</i> — Specifies the period of time before trap-suppression can occur, depending upon the number of traps received in that period of time. Values 1 to 300

tunnel-interface

Syntax [no] **tunnel-interface** **rsvp-p2mp** *lsp-name* [**sender** *sender-address*]

Context	config>router config>router>igmp
Description	This command creates a tunnel interface associated with an RSVP P2MP LSP. IPv4 multicast packets are forwarded over the P2MP LSP at the ingress LER based on a static join configuration of the multicast group against the tunnel interface associated with the originating P2MP LSP. At the egress LER, packets of a multicast group are received from the P2MP LSP via a static assignment of the specific <S,G> to the tunnel interface associated with a terminating LSP. At ingress LER, the tunnel interface identifier consists of a string of characters representing the LSP name for the RSVP P2MP LSP. The user can create one or more tunnel interfaces and associate each to a different RSVP P2MP LSP. At egress LER, the tunnel interface identifier consists of a couple of string of characters representing the LSP name for the RSVP P2MP LSP followed by the system address of the ingress LER. The LSP name must correspond to a P2MP LSP name configured by the user at the ingress LER. The LSP name string must not contain “:.” (two :s) nor contain a “.” (single “.”) at the end of the LSP name. However, a “.” (single “.”) can appear anywhere in the string except at the end of the name.
Default	none
Parameters	rsvp-p2mp <i>lsp-name</i> — Specifies the LSP. The LSP name can be up to 32 characters long and must be unique. sender <i>sender-address</i> — Specifies system address of the ingress LER for the P2MP RSVP LSP.

update-path

Syntax	update-path { lsp <i>lsp-name</i> path <i>current-path-name</i> new-path <i>new-path-name</i> }
Context	tools>perform>router>mpls
Description	This command enables you to instruct MPLS to replace the path of a primary or secondary LSP. The primary or secondary LSP path is indirectly identified via the <i>current-path-name</i> value. The same path name cannot be used more than once in a given LSP name. This command applies to both CSPF LSP and to a non-CSPF LSP. This command will only work when the specified <i>current-path-name</i> has the adaptive option enabled. The adaptive option can be enabled at the LSP level or the path level. The new path must have been configured in the CLI or provided via SNMP. The CLI command for entering the path is configure router mpls path <i>path-name</i> The command fails if any of the following conditions exist:

- The specified *current-path-name* of this LSP does not have the adaptive option enabled.

- The specified *new-path-name* value does not correspond to a previously defined path.
- The specified *new-path-name* value exists but is being used by any path of the same LSP, including this one.

When you execute this command, MPLS performs the following procedures:

- MPLS performs a single MBB attempt to move the LSP path to the new path.
- If the MBB is successful, MPLS updates the new path
 - MPLS writes the corresponding NHLFE in the data path if this path is the current backup path for the primary.
 - If the current path is the active LSP path, it will update the path, write the new NHLFE in the data path that will cause traffic to switch to the new path.
- If the MBB is not successful, the path retains its current value.
- The update-path MBB has the same priority as the manual re-signal MBB.

2.21.2.5 Clear Commands

session

Syntax	session src-ip <i>ip-address</i> dst-ip <i>ip-address</i> session src-ip <i>ip-address</i> dst-ip <i>ip-address</i> lsp-rsvp { head tail } tunnel-id [0..4294967295] lsp-id [0..4294967295] session mpls-tp <i>lsp-name</i> path { working protect } session p2mp-interface <i>interface-name</i> session src-ip <i>ip-address</i> lsp-rsvp { head tail } rsvp-session-name [256 chars max]
Context	clear>router>bfd
Description	This command clears BFD session information. Clearing the BFD session will cause it to go down and restart. This may cause any client protocols whose state is affected by BFD to go down. As in the current implementation if clear router router-instance bfd statistics all is executed, then the router-instance is ignored and the clear is applied to all session statistics.
Parameters	src-ip <i>ip-address</i> — Clears information about the specified source IPv4 or IPv6 address. dst-ip <i>ip-address</i> — Clears information about the specified destination IPv4 or IPv6 address. lsp-rsvp {head tail} — Clears information about the specified link type. tunnel-id [0..4294967295] — Clears information about the specified tunnel ID. lsp-id [0..4294967295] — Clears information about the specified LSP ID. mpls-tp <i>lsp-name</i> — Clears information about the specified MPLS TP LSP.

path {working | protect} — Clears information about the working or protect path.

p2mp-interface *interface-name* — Clears information about the specified P2MP interface.

rsvp-session-name [256 chars max] — Clears information about the specified RSVP session.

statistics

Syntax	statistics src-ip <i>ip-address</i> dst-ip <i>ip-address</i> statistics all statistics src-ip <i>ip-address</i> dst-ip <i>ip-address</i> lsp-rsvp { head tail } tunnel-id [0..4294967295] lsp-id [0..4294967295] statistics mpls-tp <i>lsp-name</i> path { working protect } statistics p2mp-interface <i>interface-name</i> statistics src-ip <i>ip-address</i> lsp-rsvp { head tail } rsvp-session-name [256 chars max]
Context	clear>router>bfd
Description	This command clears BFD statistics.
Parameters	src-ip <i>ip-address</i> — Clears statistics about the specified source IPv4 or IPv6 address. dst-ip <i>ip-address</i> — Clears statistics about the specified destination IPv4 or IPv6 address. all — Clears all statistics for the BFD instance. lsp-rsvp { head tail } — Clears statistics about the specified link type. tunnel-id [0..4294967295] — Clears statistics about the specified tunnel ID. lsp-id [0..4294967295] — Clears statistics about the specified LSP ID. mpls-tp <i>lsp-name</i> — Clears statistics about the specified MPLS TP LSP. path { working protect } — Clears statistics about the working or protect path. p2mp-interface <i>interface-name</i> — Clears statistics about the specified P2MP interface. rsvp-session-name [256 chars max] — Clears statistics about the specified RSVP session.

interface

Syntax	interface <i>ip-int-name</i> [statistics]
Context	clear>router>mpls
Description	This command resets or clears statistics for MPLS interfaces.

Parameters *ip-int-name* — The name of an existing IP interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

statistics — This parameter clears only statistics.

lsp

Syntax **lsp** *lsp-name*

Context clear>router>mpls

Description This command resets and restarts an LSP.

Parameters *lsp-name* — The name of the LSP to clear up to 64 characters in length.

lsp-autobandwidth

Syntax **lsp-autobandwidth** [*lsp-name*]

Context clear>router>mpls>lsp

Description This command clears the following counters/timers, as follows:

- The sample count is reset to zero, and the average data rate of the current sample interval is discarded.
- The adjust count is reset to zero.
- The maximum average data rate is zeroed.
- The overflow count is zeroed.

ingress-stats

Syntax **ingress-statistics**

Context clear>router>mpls

Description This command provides the context for the user to enter the LSP names for the purpose of enabling ingress data path statistics at the terminating node of the LSP (for example, egress LER).

Default none

lsp-egress-stats

Syntax	lsp-egress-stats lsp-egress-stats <i>lsp-name</i>
Context	clear>router>mpls
Description	This command clears MPLS LSP egress statistics information.

lsp-ingress-stats

Syntax	lsp-ingress-stats lsp-ingress-stats <i>ip-address</i> lsp <i>lsp-name</i> lsp-ingress-stats <i>sender-address:lsp-name</i>
Context	clear>router>mpls
Description	This command clears MPLS LSP ingress statistics information.

interface

Syntax	interface <i>ip-int-name</i> statistics
Context	clear>router>rsvp
Description	This command resets or clears statistics for an RSVP interface.
Parameters	<i>ip-int-name</i> — The name of the IP interface to clear. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. statistics — This parameter clears only statistics.

statistics

Syntax	statistics
Context	clear>router>rsvp
Description	This command clears global statistics for the RSVP instance, for example, clears path and resv timeout counters.

2.21.2.6 Debug Commands

mpls

Syntax	mpls [lsp <i>lsp-name</i>] [sender <i>source-address</i>] [endpoint <i>endpoint-address</i>] [tunnel-id <i>tunnel-id</i>] [lsp-id <i>lsp-id</i>] [interface <i>ip-int-name</i>] no mpls
Context	debug>router
Description	This command enables and configures debugging for MPLS.
Parameters	<i>lsp lsp-name</i> — Name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique. <i>sender source-address</i> — The system IP address of the sender. <i>endpoint endpoint-address</i> — The far-end system IP address. <i>tunnel-id tunnel-id</i> — The MPLS SDP ID. Values 0 to 4294967295 <i>lsp-id lsp-id</i> — The LSP ID. Values 1 to 65535 <i>interface ip-int-name</i> — Name that identifies the interface. The interface name can be up to 32 characters long and must be unique. If the string contains special characters (#, \$, spaces, etc.), 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 re-route events. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about re-route events.

iom

Syntax	iom [detail] no iom
Context	debug>router>mpls>event
Description	This command reports MPLS debug events originating from the XMA. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about MPLS events originating from the XMA.

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	[lsp <i>lsp-name</i>] [sender <i>source-address</i>] [endpoint <i>endpoint-address</i>] [tunnel-id <i>tunnel-id</i>] [lsp-id <i>lsp-id</i>] [interface <i>ip-int-name</i>] no rsvp

Context	debug>router
Description	This command enables and configures debugging for RSVP.
Parameters	<i>lsp lsp-name</i> — Name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique. <i>sender source-address</i> — The system IP address of the sender. <i>endpoint endpoint-address</i> — The far-end system IP address. <i>tunnel-id tunnel-id</i> — The RSVP tunnel ID. Values 0 to 4294967295 <i>lsp-id lsp-id</i> — The LSP ID. Values 1 to 65535 <i>interface ip-int-name</i> — The interface name. The interface name can be up to 32 characters long and must be unique. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

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 reservation events. The no form of the command disables the debugging.
Parameters	detail — Displays detailed information about RSVP reservation events.

te-threshold-update

Syntax	te-threshold-update no te-threshold-update
Context	debug>router>rsvp>event debug>router>rsvp>ip-int-name>event
Description	This command debugs the te-threshold-update events. The no form of this command disables the debugging

packet

Syntax	[no] packet
Context	debug>router>rsvp>
Description	This command enters the syntax to debug 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 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 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 resv packets.

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

Parameters **detail** — Displays detailed information about RSVP Resv events.

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.

3 GMPLS

3.1 In This Chapter

This chapter provides information to configure GMPLS.

- [GMPLS](#)
 - [Example Applications](#)
- [GMPLS UNI Architecture](#)
 - [Addressing and End-to-End gLSP Architecture](#)
- [1830 PSS Identifiers](#)
- [Recovery Reference Models](#)
 - [End to End Recovery \(IP-layer\)](#)
 - [End to End ECMP](#)
 - [End to End Load Sharing Using a Load Sharing GMPLS Tunnel Group](#)
 - [End to End Recovery \(GMPLS Layer\)](#)

3.2 GMPLS

The Generalized Multi-Protocol Label Switching (GMPLS) User to Network Interface (UNI) permits dynamic provisioning of optical transport connections between IP routers and optical network elements in order to reduce the operational time and administrative overhead required to provision new connectivity. The optical transport connections typically originate and terminate in an IP/MPLS controlled domain and traverse an intermediate optical transport network. The GMPLS UNI model is based on an overlay approach, whereby the IP/MPLS control plane is transported transparently over the intermediate transport network, which itself is controlled by a GMPLS control plane.

The UNI provides a clear demarcation point between the domains of responsibility of the parties involved in managing the overlying IP/MPLS network and the underlying optical network. For example, these parties could be two divisions in a service provider organization, or a subscriber/client of the service provider and the service provider itself.

The UNI has a client part, the UNI-C, and a network part, the UNI-N. In the Nokia solution, the UNI-C is an SR OS system, such as a 7750 SR or a 7950 XRS, while the UNI-N is an optical device; for example, an 1830 PSS.

Control plane related information is exchanged between the UNI-C and the UNI-N using a dedicated out of band communication channel. Note that the adjacent optical network element and the router assume that they are connected to a trusted peer, and thus assume a secure communication. This is achieved by physically securing the link carrying the control channel between the two.

Based on standardized UNI messaging (RFC 4208), the UNI-C indicates to the UNI-N which far-end peer UNI-C node (corresponding to a remote router) to make an optical transport connection to. This path request can include additional path attributes to indicate requirements such as bandwidth, priority and diversity/resiliency parameters.

3.2.1 Example Applications

This section summarizes some of the use cases that the GMPLS UNI may be used to address.

3.2.1.1 Use Case 1: Dynamic Connection Setup with Constraints

This use case aims to solve inefficiencies between IP and transport teams within an operator for connectivity setup; for example:

- Process complexity, with complex database exchange, parsing and filtering
- Long-winded organizational communication prior to path establishment


It therefore aims to optimize IP/Optical transport team interactions by removing complex processes, and reduces per-connection provisioning in the optical core.

The UNI should allow the setup/maintenance/release of connections across an intermediate optical transport network from a UNI-C router to another remote UNI-C router. The routers are connected to an optical network that consists of optical cross connects (OXCs), and the interconnection between the OXC and the router is based on the GMPLS UNI (RFC 4208). The UNI-C routers are 7450 ESS, 7750 SR, or 7950 XRS nodes, while the UNI-N OXC is the Nokia1830 PSS. The UNI connection is instantiated using a GMPLS LSP (gLSP).

The UNI-C router is always the initiator of the connection. The only per-connection configuration occurs at the UNI-C, and it is operator initiated. Connections to any of the remote UNI-C routers are signaled over the UNI. The initiation of a connection request is via CLI or SNMP to the UNI-C router.

Signaling is based on RSVP-TE (RFC 4208). Constraints can be signaled with a connection setup request. These include bandwidth, protection type, and latency. In the event that a connection could not be established, a correct (descriptive) error code is returned to the initiator.

Figure 40 Dynamic Connection Setup



24850


3.2.1.2 Use Case 2: Multi-Layer Resiliency

The objective of this application is to ensure optical network path diversity for primary/backup paths of an overlay IP network. It thus aims to resolve situations where the UNI-C router has no topological visibility of the optical network, and to allow the router to indicate that paths have to be either co-routed or avoid specific optical nodes or links along a path.

Route diversity for LSPs from single homed UNI-C router and dual-homed UNI-C router is a common requirement in optical transport networks. Dual homing is typically used to avoid a single point of failure (for example, the UNI link or OXC) or to allow two disjoint connections to form a protection group.

For the dual-homing case, it is possible to establish two connections from the source router to the same destination router where one connection is using one UNI link to, for example, OXC1 and the other connection is using the UNI link to OXC2. In order to avoid single points of failure within the optical network, it is necessary to also ensure path (gLSP) diversity within the provider network in order to achieve end-to-end diversity for the two gLSPs between the two routers.

Figure 41 Multi-Layer Resiliency




As the two connections are entering the provider network at different OXC devices, the OXC device that receives the connection request for the second connection needs to be capable of determining the additional path computation constraints such that the path of the second LSP is disjoint with respect to the already established first connection entering the network at a different PE device.

3.3 GMPLS UNI Architecture

This section specifies the architectural and functional elements of the GMPLS UNI on the 7450 ESS, 7750 SR, or 7950 XRS nodes and the 1830 PSS node (which must be GMRE), and how they relate to one another. The architecture is illustrated in [Figure 42](#).

Figure 42 GMPLS UNI Architecture



On the UNI-C side, the UNI consists of the following functional components:

- A set of one or more data bearers between the UNI-C and UNI-N. Each data bearer maps to a black and white Ethernet network port.
- A Traffic Engineering (TE) link (RFC 4202), represented by an identifier on the UNI-C and UNI-N nodes. This identifier is manually configured in Release 13.0. A TE link maps to a single data bearer. There may be one or more TE links per UNI between a UNI-C and UNI-N pair.
- An IP Control Channel (IPCC) between the UNI-C and UNI-N. This carries GMPLS control plane traffic between the two nodes and is separate from the links carrying user plane traffic. This carries the following two control protocols:
 - LMP — Link Management Protocol. This is responsible for checking the correlation between the UNI-C/UNI-N and the TE link/Data Bearer identifiers, and maintaining the IPCC adjacency between the UNI-C and UNI-N. LMP runs on a network IP interface bound to an Ethernet port on an Ethernet MDA/IMM. This is a separate port to the TE Links.
 - RSVP-TE — RSVP-TE runs on the same network interface as LMP. The next hop from an RSVP-TE perspective is the UNI-N. RSVP-TE is used to establish and maintain a GMPLS LSP.

- **gLSP** — The GMPLS LSP. At the UNI-C, this is a control plane object representing the TE-Link in the RSVP-TE control plane. Although this is an LSP, there is no explicit MPLS label in the data path at the UNI-C; the gLSP maps to a data bearer of the TE link to / from the UNI-N. When a gLSP is signaled to a far-end UNI-C node, the optical network establishes bidirectional connectivity between one of the data bearers in the TE link on the UNI-N at the ingress to the optical network, and one of the data bearers on the TE link on the egress UNI-N node connected to the far end UNI-C node.
- **Network Interface** — When a gLSP is successfully established, a network interface can be bound to the gLSP. The network interface then uses the data bearer associated with the gLSP to forward traffic. This network interface can be used by any applicable protocol associated with an overlying IP/MPLS network. The network interface is bound to the gLSPs via a GMPLS tunnel group.
- **GMPLS Tunnel Group**: A GMPLS tunnel group is a bundle of gLSPs providing an abstraction of the data bearers that are intended to be associated to one IP interface. A GMPLS tunnel group only exists on the UNI-C node and not on the 1830 PSS UNI-N node.

Although [Figure 42](#) shows a single 7450 ESS, 7750 SR or 7950 XRS node connected to a single UNI-N (1830 PSS), it is possible to multi-home a router into more than one (usually two) UNI-Ns at the edge of the optical network. In this case, a separate IPCC, set of data bearers, and set of TE links, are required between the 7450 ESS, 7750 SR or 7950 XRS and each UNI-N.

3.3.1 Addressing and End-to-End gLSP Architecture

The GMPLS UNI assumed a flat addressing scheme between the UNI-C nodes and the optical network. In this model, a common addressing scheme is used between the UNI-C (IP router) and UNI-N (optical edge). The UNI-C and UNI-N must be in the same subnet. Also, none of the UNI-C addresses can overlap or clash with any of the GMPLS-aware nodes in the optical network. This does not mandate that the whole IP network share a common address space with the optical network, as a separate loopback address can be used for the GMPLS UNI on the UNI-C.


The ERO Expansion (RFC 5151) model is assumed for the GMPLS LSPs. The UNI-C is not exposed to the full ERO between the UNI-N nodes. Instead, the full ERO is inserted at the UNI-N. This model limits the sharing of topology information between the UNI-N and UNI-C.

3.4 1830 PSS Identifiers

This section describes the various identifiers used on the 1830 PSS that are relevant to configuring the GMPLS UNI on the 7450 ESS, 7750 SR, or 7950 XRS node in conjunction with the 1830 PSS. [Figure 43](#) illustrates the identifier architecture of a 1830 PSS multi-shelf system. The multi-shelf system consists of a control plane node and one or more data plane nodes. The following identifiers are used:

- GMRE node IP— This is the IP loopback address used for GMPLS protocols such as LMP and RSVP.
- IPCC IP address (also known as DcnGatewayAddress)— This is the source/destination address for IPCC maintenance messages such as LMP hellos and LMP config messages. When only one IPCC exists between a 7450 ESS, 7750 SR, or 7950 XRS and 1830 PSS pair, this may be the same as the IP management loopback.
- CP Node ID — This is a non-routable identifier for the control plane node. It is used for identifying this node in the optical domain; for example, the session/sender template. It is also used in the RSVP ERO to identify the 1830 PSS node.
- DP Node ID — This is a non-routable identifier for a data plane node. This identifies a particular data plane shelf in the optical domain.
- TE Link IDs — The TE Link ID is unique across a set of DP and CP nodes forming an 1830 PSS.

Figure 43 Identifier Architecture




3.5 Recovery Reference Models

This section details the supported recovery reference models. These models are based on the mechanisms specified in RFC 4872 and RFC 4873.

Figure 44 presents a generalized reference model in which the 7450 ESS, 7750 SR, or 7950 XRS UNI-C nodes are dual-homed at the link layer to the 1830 PSS UNI-N nodes. Not all elements of this architecture may be required in all deployment cases.

Figure 44 General GMPLS UNI Interconnection Architecture



This reference model includes two 7450 ESS, 7750 SR, or 7950 XRS nodes, each hosting a UNI-C function, at the edge of each IP network facing two 1830 PSS nodes, each hosting a UNI-N function. A full mesh of black and white Ethernet links interconnects neighboring UNI-C nodes and UNI-N nodes. Parallel links may exist, so that a given 7450 ESS, 7750 SR, or 7950 XRS UNI-C is connected to a neighbor 1830 PSS UNI-N by more than one Ethernet link.

Each router hosting a UNI-C has an IPCC to each of the two 1830 PSS UNI-Ns. Likewise, each 1830 PSS hosting UNI-N has an IPCC to both of the 7450 ESS, 7750 SR, or 7950 XRS UNI-Cs that it is connected to. IPCCs only exist between UNI-C and UNI-N nodes, and not between UNI-C nodes. A control plane (LMP and RSVP) adjacency therefore exists between each UNI-C and its corresponding UNI-Ns.

Recovery in the following domains is supported in the following locations:

- End to End — Between the 7450 ESS, 7750 SR, or 7950 XRS UNI-C nodes at each end of a gLSP.
- Optical Segment — Between 1830 PSS UNI-N nodes at each edge of the optical network.

The following subsections detail some example recovery options that are possible using either GMPLS, or a combination of GMPLS mechanisms and mechanisms in the overlay IP network. Note that some of the functionality shown in one of the scenarios can be used in combination with functionality in another scenario, for example optical SRLG diversity.


The objective of GMPLS here is to minimize the disruption to the overlay IP network while simultaneously maximizing the utilization of both the gLSPs and the resources in the underlying optical network (or UNI links).

3.5.1 End to End Recovery (IP-layer)

End to end recovery applies to protection against failures at any point along the entire path between a local UNI-C and a far end UNI-C. In the context of the GMPLS UNI, recovery can be implemented in the overlay IP network either at Layer 3 or Layer 2, with assistance from the underlay optical network, with optional additional protection and/or restoration of gLSPs by GMPLS.

3.5.2 End to End ECMP

[Figure 45](#) illustrates the first model. Multiple gLSPs are established between a pair of remote UNI-C nodes. Each gLSP is bound to a separate IP network interface at the UNI-C. RSVP signaling across the UNI is used to ensure that the gLSPs are SRLG diverse (by explicitly signaling the SRLG list to avoid in an XRO for every gLSP, or automatically collecting the SRLG list for a gLSP which does not have an XRO, and then signaling a subsequent gLSP including this collected list in its XRO). Protection is provided at the IP layer by hashing across the IP network interfaces associated with each gLSP. The operational state of each IP interface can be tied to the operational state of its gLSP (controlled using RSVP) or using mechanisms in the IP overlay such as BFD.


Figure 45 End-to-End ECMP with gLSP Diversity Across Single UNI-C

24856

3.5.3 End to End Load Sharing Using a Load Sharing GMPLS Tunnel Group

Figure 46 shows the case where multiple gLSPs, instantiated as black and white Ethernet ports, are bundled together in a similar manner to LAG, using a GMPLS tunnel group. That is, each member gLSP of a tunnel group effectively maps to a member port, which runs end to end between remote UNI-Cs. Note that a LAG does not and cannot terminate on the neighboring 1830 PSS UNI-N. A single IP network interface is bound to the bundle of ports represented by the gLSPs. LACP does not run across the bundle; RSVP signaling is instead used to convey the state of the gLSP and thus the corresponding member port of the tunnel group. Traffic is load shared across the tunnel group members.

Figure 46 End-to-End Load Sharing GMPLS Tunnel Group with gLSP Path Diversity




24857

3.5.4 End to End Recovery (GMPLS Layer)

3.5.4.1 Unprotected gLSP

The default level of E2E recovery is unprotected. In this case, a gLSP can only recover from a failure when the downstream resource that failed is recovered. [Figure 47](#) illustrates this. When a gLSP fails in the optical network, a failure notification is propagated to the source UNI-C node e.g. using a PathErr or a NotifyErr LSP Failure message. The source UNI-C node takes no action, but will continue to refresh the PATH message for this gLSP, which may be rerouted around the failure by the optical network e.g. if the IGP in the optical network reconverges. The gLSP is treated as operationally down until a message indicating that the gLSP has been restored is received by the router; for example, a Notify Error LSP Restored.


Figure 47 gLSP Re-Establishment (PATH Refresh)

24858

3.5.4.2 Full LSP Rerouting

Full LSP rerouting (or restoration), specified in RFC 4872 section 11, switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node; it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links.

Figure 48 Full LSP Rerouting




24859

3.5.4.3 1: N Protection


In 1:N ($N \geq 1$) protection, the protecting LSP path is a fully provisioned and resource-disjoint LSP path from the N working LSP paths. The N working LSP paths may also be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working paths to the protecting path. Although RFC4872 allows extra traffic on the protecting path, this is not supported by the 7450 ESS, 7750 SR, or 7950 XRS. [Figure 49](#) illustrates this protection architecture when $N=1$, while [Figure 50](#) shows the case for $N>1$.

Figure 49 **1:N Protection, with N=1 (RFC4872)**



al_0712

Figure 50 **N>1 Protection**




al_0906

3.5.4.4 Optical Segment Recovery

Optical segment protection refers to the ability of the optical network to protect the span of a gLSP between ingress and egress UNI-N nodes. It does not require any protection switching on the UNI-C nodes. However, it does require the UNI-C to signal a request for a particular segment protection type towards the UNI-N in the PATH message for a gLSP. The optical network may either accept this request, reject it or respond with an alternative. Segment protection is defined in RFC 4873.

Figure 51 Optical Segment Protection Domain



24860

Signaling of the following segment protection types is supported by the 7450 ESS, 7750 SR, and 7950 XRS:

- **Unprotected** — The path is not protected against failure.
- **Source-Based Reroute (SBR)** — In this mechanism (also known as Full Rerouting), a path is restored after a failure, but the success of restoration depends on the available resources. This can reroute traffic in 200 ms or more.
- **Guaranteed Restoration (GR)** — A shared backup is assigned to the path, and recovery resources are reserved. If they cannot be reserved on a shared path, then this falls back to SBR. This can reroute traffic in 50 ms or less. This mechanism is also known as 1+shared standby. This is also known as Rerouting without extra traffic, or shared mesh restoration.
- **Sub-network Connection Protection (SNCP)** — This provides 50 ms protection in the case of a single failure. This is also known as 1+1 bidirectional path protection.

- Path Restoration Combined (PRC) — This provides 50 ms protection, even in the case of multiple failures. This is also known as SNCP with SBR.

3.6 Configuring GMPLS with CLI

This section provides information to configure UNI GMPLS using the command line interface.

Topics in this section include:

- [GMPLS Configuration Overview](#)
- [LMP and IPCC Configuration](#)
 - [Configuration of IP Communication Channels for LMP and RSVP](#)
 - [Configuring LMP](#)
 - [Configuring Traffic Engineering Links and Data Bearers](#)
- [Configuring MPLS Paths for GMPLS](#)
- [Configuring RSVP in GMPLS](#)
- [Configuring a GMPLS LSP on the UNI](#)
 - [gLSP Constraints](#)
- [Bandwidth](#)
- [Shared Risk Link Groups](#)
- [Optical Network Segment Recovery](#)
- [Configuration of End-to-End GMPLS Recovery](#)
- [GMPLS Tunnel Groups](#)
- [Configuring IP and MPLS in an Overlay Network to Use a GMPLS LSP](#)
- [Configuration Notes](#)

3.7 GMPLS Configuration Overview

The Generalized Multi-Protocol Label Switching (GMPLS) User to Network Interface (UNI) permits dynamic provisioning of optical transport connections between IP routers and optical network elements in order to reduce the operational time and administrative overhead required to provision new connectivity.

3.8 LMP and IPCC Configuration

3.8.1 Configuration of IP Communication Channels for LMP and RSVP

Configuration starts with enabling the IP communication channel (IPCC) between the 7450 ESS, 7750 SR, or 7950 XRS UNI-C and the adjacent UNI-N. The IPCC is a data communication channel for LMP and RSVP. For each different UNI-C and UNI-N adjacency, a different IPCC must be configured.

In Release 13.0, a numbered network IP interface is bound to the port connected to the DCN or directly to the 1830 PSS.

GMPLS protocols use a new loopback address type, called a **gmpls-loopback**, on the IPCC. The address of this loopback is termed the local GMPLS router ID. Packets that do not belong to a GMPLS protocol that are destined for this loopback address will be dropped. An interface is configured as a GMPLS Loopback using the **gmpls-loopback** keyword.

```
config
router
  interface local-gmpls-router-id-name
    gmpls-loopback
    address local-gmpls-loopback-address //Local LmpNodeId
```

The destination address of the LMP and RSVP control plane packets should be set to the LMP/GMPLS loopback of the 1830 PSS. The 1830 PSS does that via a dedicated subnet on a VLAN interface on the management port. Another VLAN extends a separate subnet for management purposes. On the 7450 ESS, 7750 SR, or 7950 XRS, the LMP and RSVP control plane packets should be sent to the next-hop for the GMPLS/LMP loopback address of the neighboring 1830 PSS. This is achieved via a static route in Release 13.0. The 1830 PSS and the GMPLS Router IDs must be in the same subnet. It may be possible to operate over a routed DCN network if the RSVP control plane messages will not set the IP router alert bit. Otherwise only direct IP connectivity, via a L2 network, will work.

If the IPCC goes down, then an existing TE Link or gLSP to a given peer UNI-N node is not torn down just because the IPCC is down. However, if the IPCC is down, then it is not possible to establish new gLSPs or TE Links, and a trap indicating a degraded state is raised.

3.8.2 Configuring LMP

LMP is used to establish and maintain an IPCC between adjacent peers, as well as to correlate the local and remote identifiers for the TE Links that it controls. Some attributes must be configured locally on a per-peer basis, such as the LMP peer information, te-link information, and per-peer protocol related parameters.

The **config>router>imp>imp-peer** *peer-cp-node-id* command creates a context per LMP peer. The entry *peer-cp-node-id* is the control plane identifier of the adjacent UNI-N. It is an IPv4 or unsigned integer-formatted address that is used by the UNI-C for LMP and RSVP-TE messages if a peer-loopback address is not subsequently configured. The local GMPLS Router ID is used as the source address.

In Release 13.0, a static route must have previously been configured to this peer router ID. Dynamic routing e.g. using OSPF over the IPCC in order to resolve routes to the peer GMPLS router ID, is not supported. In addition, the local loopback address to use as the local GMPLS Router ID should also be configured.

The LMP messages are sent over the interface corresponding to the IPCC that has been configured previously. The LMP session can be associated with one or more TE links that have been configured previously.

A control channel to an LMP Peer is configured using the **config>router>imp>imp-peer** *peer-cp-node-id>control-channel* context. Control channels are indexed using the *imp-cc-id* parameter, which corresponds to the *ImpCcid* object in the LMP MIB.

The following CLI tree illustrates the key commands for configuring LMP.

```
config
  router
    [no] imp
      [no] te-link te-link-id
      link-name te-link-name
      remote-id id
      [no] data-bearer data-bearer-id
      port port-id
      remote-id id
      [no] shutdown
      [no] shutdown

    gmpls-loopback-address local-gmpls-loopback-address
    [no] imp-peer peer-cp-node-id
      peer-loopback-address peer-loopback-address
      retransmission-interval interval
      retry-limit limit
      [no] control-channel imp-cc-id
        peer-interface-address ipcc-destination-addr
        hello interval interval dead-interval interval
        passive
        [no] shutdown
      te-link te-link-id
```

```
te-link te-link-id
[no] shutdown
  lmp-peer lmp-peer-address
  ...
  [no] shutdown
[no] shutdown
```

If **peer-loopback-address** is entered, then this is used as the routable peer address, otherwise the *peer-cp-node-id* is assumed to correspond to a routable peer loopback.

The **peer-interface-address** is mandatory and is the destination address of the IPCC on the peer UNI-N used to reach the GMPLS Router ID of the peer. It corresponds to the `ImpCcRemotelpAddr` in RFC 4631. If the **peer-interface-address** is used as the destination IP address in the IP packet on the IPCC, then the router local interface address is used as the source IP address.

A **te-link** is configured under **config>router>lmp>te-link**. The **te-link** parameter under **config>router>lmp>lmp-peer** then assigns the control of the TE-links to the LMP protocol to a given peer. Each TE-Link can only be assigned to a single LMP peer.

The LMP protocol-specific attributes such as timers and retransmission retries are configured for each LMP peer under **configure>router>lmp>lmp-peer**.

The **hello interval** ranges from 1000 to 65 535 ms. The default hello interval is 1000 ms.

The **hello dead-interval** ranges from 3000 to 65 535 ms. The default hello dead interval is 4000 ms.

The **retransmission-interval** ranges from 10 to 4 294 967 295 ms in 10-ms intervals, with a default of 500 ms.

The **ttl** command allows the user to configure the TTL of the IP control channel for RSVP and LMP packets to a value other than 1 (default). The range of values is 2 - 255. This enables multi-hop data communication networks between the UNI-C and UNI-N.

3.8.3 Configuring Traffic Engineering Links and Data Bearers

Traffic engineering (TE) links are configured under the **config>router>imp** with a specific command, **te-link**, to create a specific context to hold TE specific configuration information pertinent to the local and remote identifiers, and physical resources assigned to the te-link. Only one data bearer per TE link is supported.

The te-link association is the creation of an association between a TE-link and data-bearing physical ports. Under the TE-link context, different data bearers can be configured via the data-bearer command. The data bearer is assigned a complete physical port, using `port<x/y/z>` (slot-number/MDA-number/port-number) as input.

Note that a data bearer cannot be associated with a port in a LAG.

A TE-link has a unique *link-id*, which identifies it in RSVP-TE signaling.

The remote-id is the unnumbered link identifier at far-end of the TE link as advertised by the LMP peer i.e. the UNI-N.

The TE-link has associated physical resources which are assigned to the TE-link by configuring the data-bearer under the **config>router>te-link** context.

The operator must also configure the remote data-bearer link identifier under the data bearer subcontext.

Note that LMP does not correlate the local and remote Layer 2 interface identifiers (such as MAC addresses) for the data bearer. It only correlates the local and remote TE Link and Data Bearer link identifiers. The association between the Layer 2 interface address and the data bearer must be correctly configured at the UNI-C and UNI-N. The **show>router>imp>te-link** command displays the local link ID, remote link ID, and associated port ID to assist with this.

The CLI tree for creating TE Links under LMP is as follows. Note that there are also some RSVP-specific TE Link parameters that are configured under a separate **gmpls** context (see below):

```
config
  router
    [no] lmp
      [no] te-link te-link-id
        link-name te-link-name
        remote-id id
      [no] data-bearer data-bearer-id
        port port-id
        remote-id id
      [no] shutdown
```

```
[no] shutdown
[no] shutdown
```

The *te-link-id* can take the form of an unsigned integer or 64 character (max) name: [1..2147483690] | *te-link-name*: 64 char max

Upon creation, only the unsigned integer needs to be specified. Once the link is created the user can configure the link name (ie. 'link-name *te-link-name*'). From here, the user can refer to this te-link by either the unsigned integer or the ASCII name.

Note that LMP will normally assume a data bearer is operationally up, even if no MAC layer or a valid PCS IDLE stream is received. This is because a neighboring UNI-N may not generate a PCS IDLE stream and instead transparently transports the MAC layer from the far end, which won't be up unless a gLSP is configured. In order to prevent LMP from using a port for which there is a local fault on the data bearer, indicated by loss of light, a user must configure **report-alarm** on the Ethernet port, as follows:

```
config>port>ethernet>report-alarm signal-fail
```

Only ports with **report-alarm signal-fail** configured can be included in LMP, and that **report-alarm signal-fail** cannot be subsequently removed from a port in LMP.

RSVP requires that all traffic engineering attributes for TE Links are configured under the **config>router>gmpls>te-link** context.

```
config
router
  [no] gmpls
    te-link te-link-id
      [no] shutdown
```

where *te-link-id*: [1..2147483690] | *te-link-name*: 32 char max

If a path (also refer to the description of a GMPLS path configuration, below) without an explicit te-link for the first hop is configured, the system will automatically select a TE Link to use for a gLSP path based on the lowest available TE Link ID with a matching bandwidth (if a bandwidth is configured for the gLSP). During a data-bearer link allocation request, an RSVP -requested gLSP BW could be either a non-zero value as per RFC 3471 signal-type (see below), or it could be zero. These are the following cases.

Case 1: Requested BW is non-zero as per RFC 3471 Signal-type configuration

- When a TE (or TE/DB) link is configured in the related hop LMP checks whether the related port BW is the same (exact match) as the requested BW, and allocates the port (provided any other checks are successful).

- When the related Hop is empty, LMP finds a db-link port to the peer with a matching the requested BW, and allocates it.

Case 2: Requested BW is Zero

- When TE (or TE/DB) link is configured in the related hop, LMP allocates the port (provided the other checks are OK), and provides the port BW to RSVP to use in signaling.
- When the related Hop is empty, LMP finds the first available db-link to the peer (based on lower db-link Id), and allocates it and provides the port BW to RSVP to use in signaling.

3.9 Configuring MPLS Paths for GMPLS

To establish an end-to-end connection between two 7450 ESS, 7750 SR, or 7950 XRS nodes through a GMPLS network, a path is required, which is configured via the **configure>router>gmpls>path** *path-name* context.

The path context consists of a set of numbered entries, each entry representing a resource that the gLSP must follow. The te-link ID is the ID allocated at the node referred to in the hop.

When interoperating with the Nokia 1830 PSS, at least the first and penultimate hops of the gLSP should be included.

The following CLI tree is used to configure a gLSP path:

```
config>router>gmpls
  path path-name
  no path path-name
    hop hop-index node-id node-id [te-link te-link-id]
      [strict | loose]
    no hop hop-index
    no shutdown
    shutdown
```

where:

node-id: IPv4 address a.b.c.d | 1830-data-plane-node-id 32-bit unsigned integer

In general, the 7450 ESS, 7750 SR, or 7950 XRS node is able to populate the ERO with every hop along the gLSP path from ingress UNI-N to egress UNI-C. However, normally only a loose path across the optical network (from ingress UNI-N to egress UNI-N) is required because the optical network is responsible for path selection between ingress and egress UNI-N. Therefore the user will normally just configure hop 1 and hop 4 in the above example. For interoperability with the 1830 PSS, the user must configure a TE Link ID to use on the final hop in the ERO towards the destination UNI-C.

The following example shows how the Path should be configured for interoperability with the 1830 PSS.

Consider the following topology:

A B C D E F

[unic1]-----[unin1]-----[unin2]-----[unic2]

where A-F are the TE Link IDs assigned at each end of a link.

Path configuration on unic1:

Hop 1 unic1 A strict

Hop 2 unin2 E loose

3.10 Configuring RSVP in GMPLS

RSVP-TE must be enabled on the SR OS towards the adjacent UNI-N in order to configure a GMPLS label-switched path (gLSP).

RSVP parameters specific to GMPLS are configured under the **config>router>gmpls** context.

This creates a new instance of RSVP for use in GMPLS signaling.

Global parameters for GMPLS are configured as follows:

```
config
  router
    gmpls
    no gmpls
    peer peer-cp-node-id
    gr-helper-time max-recovery recovery-interval max-restart restart-interval
    no gr-helper-time
    keep-multiplier number
    no keep-multiplier
```

```
no rapid-retransmit-time
rapid-retransmit-time hundred-milliseconds
no rapid-retry-limit
rapid-retry-limit limit
no refresh-time
refresh-time seconds
no refresh-time
lsp-init-retry-timeout seconds
no lsp-init-retry-timeout
no shutdown
shutdown
```

The default max-restart interval for GMPLS is 180 s.

The LMP Peer is configured under **config>router>gmpls>peer** *peer-cp-node-id*, where the *peer-cp-node-id* is control plane identifier of the adjacent optical cross connect (UNI-N node). RSVP uses the destination address returned by LMP for this peer control plane node ID as the destination address, and the loopback address referenced under **config>router>lmp>gmpls-loopback-address** *local-gmpls-loopback-address* as the local router ID to use for the session.

RSVP will come up if at least one IPCC is up.

RSVP hellos and support for graceful restart helper functionality are supported. RSVP Graceful Restart Helper procedures implemented by the router also apply when the IPCC goes down and comes back up, or when the neighboring peer control plane restarts.

The following CLI tree is used for configuring RSVP parameters for each LMP peer:

```
config
  router
    gmpls
      peer peer-cp-node-id
      no peer peer-cp-node-id
      lsp-hold-timer hold-timer
      no lsp-hold-timer
      hello-interval milliseconds
      no shutdown
      shutdown
```

The per-peer **lsp-hold-timer** *hold-timer* parameter is used to configure a node-wide hold-down time. This timer is started when a RESV for a new gLSP is first received, or a failed gLSP path is restored (or the router is notified of a restoration following segment recovery) in order to give the optical network time to program its data path. The value range is 5 to 300 s, with a default of 60 s. A member of a GMPLS tunnel group is not considered up until the hold-timer has expired. Note that different optical network technologies have different data path programing/setup times.

Note that the **no hello-interval** command sets the hello-interval to the default value of 3000 ms. Configuring **hello-interval 0** will disable hellos in GMPLS.

3.11 Configuring a GMPLS LSP on the UNI

A GMPLS LSP is configured under **config>router>gmpls>lsp name gmpls-uni**. The optional **gmpls-uni** keyword indicates that the LSP is an RSVP signaled GMPLS LSP, which is profiled for the GMPLS UNI i.e. it uses the set of functions and CLI commands applicable to an overlay gLSP, rather than a peer model gLSP. Only overlay model gLSPs are supported in Release 13.0; this is the default type of GMPLS LSP. The router can only act as an LER terminating a gLSP, and cannot switch a GMPLS i.e. it cannot act as a GMPLS LSR

GMPLS LSPs use the working path and protect path terminology from RFC 4872. Each gLSP configuration is composed of a working path and an optional protect path if end-to-end recovery is used.

Note that on-the-fly changes to an LSP or LSP path configuration are not allowed. This is because MBB is not supported for gLSPs. The LSP or LSP Path must be shut down to make configuration changes.

A GMPLS LSP (gLSP) is configured using the following CLI tree:

```
config
  router
    gmpls
      lsp lsp-name [gmpls-uni]
      no lsp lsp-name
        to remote-uni-c-gmpls-router-id
        switching-type {dcsc}
        no switching-type
        encoding-type {line}
        no encoding-type
        generalized-pid {ethernet}
        no generalized-pid
        e2e-protection-type {unprotected | 1toN | sbr}
        no e2e-protection-type
        protect-path path-name
        no protect-path path-name
          peer peer-gmpls-router-id
          no peer
          bandwidth signal-type rfc3471-name
          no bandwidth exclude-srlg group-name [group-name...(upto 5 max)]
          no exclude-srlg
          segment-protection-type {unprotected | sbr | gr | sncp | prc}
          no segment-protection-type
          no shutdown
          shutdown
        revert-timer timer-value //1 to 1800 seconds, default 0
        no revert-timer
        retry-limit limit
        no retry-limit
        no shutdown
        shutdown
        working-path path-name
        no working-path path-name
```

```

bandwidth signal-type rfc3471-name
no bandwidth
exclude-srlg group-name [group-name...(upto 5 max)]
no exclude-srlg
peer peer-gmpls-router-id
no peer bandwidth
segment-protection-type {unprotected | sbr | gr | sncp | prc}
no segment-protection-type
no shutdown
shutdown
no shutdown
shutdown

```

The loopback address of the remote router (UNI-C) must be configured after the **to** keyword and takes an IPv4 address as input.

The **switching-type** indicates the type of switching required for the gLSP. This can take a number of values, as defined in RFC 3471, and extended in RFC 6004 and RFC 7074 for Ethernet VPL (EVPL) services. The default CLI value is **DCSC**. This is the only supported value in Release 13.0.

The **encoding-type** configuration specifies the encoding type of the payload carried by the gLSP. **line**, indicating 8B/10B encoding, is the only supported type in Release 13.0.

The **generalized-pid** parameter specifies the type of payload carried by the gLSP. Standard ethertype values are used for packet and Ethernet LSPs (see RFC 3471). Only Ethernet (value 33) is supported in Release 13.0.

Note that gLSPs are inherently bidirectional. That is, both directions of the gLSP are bound together. The destination UNI-C node will automatically bind an incoming gLSP PATH message to the corresponding egress direction based on the session name in the session object.

Any gLSP that needs to be bound to a specific TE Link (as referred to in the pPATH), will only be allowed if the corresponding TE Link exists under config>router>gmpls. Constraints such as HOP definition, SRLG, BW, etc., will be checked before signaling the gLSP.

Since RSVP signaling operates out of band, refresh reduction is not supported. RSVP authentication is not supported on the 1830 PSS UNI-N, but MD5 authentication is implemented.

A configurable **retry-timer** is supported.

A configurable **retry-limit** for each gLSP is supported, with a range of 0 to 10000, and a default of 0.

The **working-path** and **protect-path** command allows paths to be configured for the gLSP. At least a **working-path** must be configured, although the path-name that it references may contain an empty path. The optional **working-path>peer** and **protect-path>peer** commands allow the user to specify a first hop UNI-N node to use for the gLSP path. The protect path is only configurable for 1:N recovery option.

Reversion from the protect path to the working path is supported.

RSVP uses the Fixed Filter (FF) style of RESV. The signaled MTU is hard-coded to 9212 bytes, as appropriate for Ethernet gLSPs.

The default **setup** and **hold** priorities are 5 and 1, respectively, and cannot be configured in Release 13.0. gLSP preemption is not supported.

Record and **record-label** are enabled by default and no user configurable command is therefore provided.

3.11.1 gLSP Constraints

Each gLSP can be configured with the following constraints:

- Bandwidth
- SRLG
- Protection

3.12 Bandwidth

The bandwidth associated with a gLSP is configured with the bandwidth command, and can take the RFC 3471 signal type name as input in Release 13.0.

The signaled bandwidth is then used for path computation and admission in the GMPLS domain.

By default, the actual interface bandwidth is used. If the user configures a bandwidth greater than the local data bearer bandwidth, then the gLSP establishment will be blocked. If the user configures a bandwidth less than or equal to the local data bearer bandwidth, then that bandwidth is signaled to the UNI-N.

The bandwidth required for the LSP is configured under the path context as follows. Note that the system will do an exact match check of the gLSP bandwidth against the data bearer bandwidth:

```

config
router
  gmpls
    lsp gmpls-tunnel-name [gmpls-uni]
      to remote-uni-c-gmpls-router-id
      working-path path-name
      bandwidth signal-type rfc3471-name

```

The possible signal-type values are:

ds0 | ds1 | e1 | ds2 | e2 | ethernet | e3 | ds3 | sts-1 | fast-ethernet | e4 | fc-0-133m |
oc-3/stm-1 | fc-0-266m | fc-0-531m | oc-12/stm-4 | gige | fc-0-1062m | oc-48/stm-16
| oc-192/stm-64 | 10gige-ieee | oc-768/stm-256 | 100gige-ieee


The code points to use for 10gige-ieee and 100gige-ieee are not yet registered with IANA. The following values are therefore used:

- 10G IEEE: 0x4E9502F9
- 100G IEEE: 0x503A43B7

3.13 Shared Risk Link Groups

Shared Risk Link Groups (SRLG) are used in the context of a gLSP to ensure that diverse paths can be taken for different gLSPs through the optical network. For example, consider the network shown in [Figure 52](#):

Figure 52 SRLG Example



24854

In this dual-homing scenario, the primary gLSP takes TE-Link 1-A, and C-2, while the secondary gLSP path takes TE-Links 1-D and F-2. In order to ensure that a failure in the underlying optical network does not affect both the primary and secondary paths for the gLSP, the SRLG list used by the optical network for the primary path is shared with the UNI-C (1) by the UNI-N (A) at the time the gLSP is established along the primary path. When the secondary path is signaled, the UNI-C (1) will signal the SRLG list to avoid to the UNI-N (D). Note that a similar procedure is beneficial even if a UNI-C is not dual homed to the optical network, but diverse primary and secondary paths are required through the optical network.

The 7450 ESS, 7750 SR, and 7950 XRS routers support two methods for indicating a set of SRLGs to exclude:

- Explicit configuration of an SRLG list for a gLSP path. These are signaled in the XRO of the RSVP PATH message towards the optical network
- Automatic SRLG collection for a gLSP, using the procedures specified in draft-ietf-ccamp-rsvp-te-srlg-collect-04.txt, and operate as follows:
 - Retrieving SRLG information from a UNI-N for an existing gLSP Path — When a dual-homed UNI-C device intends to establish a gLSP path to the same destination UNI-N device via another UNI-N node, it can request the SRLG information for an already established gLSP path by setting the SRLG information flag in the LSP attributes sub-object of the RSVP PATH message using a new SRLG flag. This path would be the primary path for a gLSP established by the router UNI-C. As long as the SRLG information flag is set in the PATH message, the UNI-N node inserts the SRLG sub-object as defined in draft-ietf-ccamp-rsvp-te-srlg-collect-04.txt into the RSVP RESV message that contains the current SRLG information for the gLSP path. Note that the provider network's policy may have been configured so as not to share SRLG information with the client network. In this case the SRLG sub-object is not inserted in the RESV message even if the SRLG information flag was set in the received PATH message. Note that the SRLG information is assumed to be always up-to-date by the UNI-C.
 - Establishment of a new gLSP path with SRLG diversity constraints — When a dual-homed UNI-C device sends an LSP setup requests to a UNI-N for a new gLSP path that is required to be SRLG diverse with respect to an existing gLSP path that is entering the optical network via another UNI-N, the UNI-C sets a new SRLG diversity flag in the LSP attributes sub-object of the PATH message that initiates the setup of this new gLSP path. This path would be the protect path of a gLSP established by the router. When the UNI-N receives this request it calculates a path to the given destination and uses the received SRLG information as path computation constraints.

In Release 13.0, the router collects SRLG by default. SRLG collection occurs on all paths of the gLSP. The collected SRLG list is visible to the user via a **show** command. The recorded SRLGs are then used to populate the XRO. Only best effort (ie. loose) SRLG diversity is supported.

Automated SRLG diversity is supported for the working and protect paths of the following end to end protection types in Release 13.0R1:

- 1:N
- LSPs that form a part of a load sharing tunnel group

Already-established gLSPs within a load-sharing tunnel group or for which 1:N recovery is configured can be made mutually diverse by applying a **shutdown / no shutdown** operation. GMPLS LSPs with other types of protection can be made mutually SRLG-diverse by performing a shutdown of the gLSP, reconfiguring the SLG list to exclude using the **exclude-srlg** command, and then applying a **no shutdown** of the gLSP.

3.14 Optical Network Segment Recovery

The router may request a particular GMPLS recovery type for a gLSP path segment that spans the optical network. This refers to the protection afforded to the gLSP path between the UNI-N nodes. The router supports the following segment protection types (code points are also shown):

- Unprotected: 0x00
- Source-Based Reroute (SBR) (Known as Full Rerouting in the IETF): 0x01
- Guaranteed Restoration (GR) (Also known as shared mesh restoration): 0x02
- Sub-network Connection Protection (SNCP) (1+1 bidirectional protection): 0x10
- Path Restoration Combined (PRC): 0x11

These resiliency options are configured under the **segment-protection-type** command for a given path.

```
config
router
  gmpls
    lsp gmpls-tunnel-name [gmpls-uni]
      to remote-uni-c-gmpls-router-id
      working-path path-name
        [no] segment-protection-type {unprotected | sbr | gr | sncp | prc}
        ...
        [no] shutdown
```

The default **segment-protection-type** setting is **unprotected**.

If the requested protection type cannot be satisfied by the optical network, the router will generate a CLI warning and an SNMP trap.

Table 29 lists the recommended combinations of segment protection type and end-to-end protection type.

Table 29 Combinations of End-to-End and Segment Protection

E2E/Segment	Unprotected	SBR	GR	SNCP	PRC
Unprotected	Yes	Yes	Yes	Yes	Yes
1:1/1:N	Yes	Yes	Yes	Yes	No
Full Rerouting	Yes	No	No	Yes	No

3.15 Configuration of End-to-End GMPLS Recovery

End-to-end GMPLS recovery is configured at the LSP level using the **e2e-protection-type** command, as follows:

```
config
router
  gmpls
    lsp gmpls-tunnel-name [gmpls-uni]
      to remote-uni-c-gmpls-router-id
      e2e-protection-type [unprotected | 1toN | sbr]
      revert-timer timer-value
```

The protection type names are common to those used in the optical network. The protection types are as follows:

- **unprotected** — 0x00
- **1toN** — 1:N protection. Extra traffic is not supported. Note that 1:1 protection is a special case of 1:N. 0x04
- **sbr** — Full LSP rerouting; 0x01

The default end-to-end protection type is **unprotected**.

It is possible to configure segment protection on a path independently of the type of end-to-end protection that is configured.

1toN protection requires the configuration of multiple working paths and a protect path for a GMPLS LSP. The working paths are then associated with different GMPLS Tunnel Groups. Configuration is as follows:

```
config
router
  gmpls
    lsp lsp-name gmpls-uni
      to remote-uni-c-gmpls-router-id
      e2e-protection-type 1toN // Only these types are allowed for gmpls-uni
      switching-type ethernet
      encoding-type ethernet
      generalized-pid ethernet
      revert-timer timer-value
      retry-limit limit
      working-path path-name1 [lmp-peer <peer-gmpls-router-id>] ...
        [no] shutdown
      working-path path-name2 [lmp-peer peer-gmpls-router-id] ...
        [no] shutdown
      working-path path-name3 [lmp-peer peer-gmpls-router-id] ...
        [no] shutdown
      protect-path path-name4 [lmp-peer peer-gmpls-router-id] ...
        [no] shutdown
```

The LSP is then bound to one or more GMPLS tunnel groups. Load sharing or 1:N protection may be used across the working paths. The load sharing case is described below.

For the non-load sharing 1:N case, a single LSP is assigned to each tunnel group, as follows:

For the head end node:

```
config > gmpls-tunnel-group 2 create
  type head-end
  far-end remote-uni-c-router-id
  mode protection
  member 1 create
    glsp session-name lsp-name:path-name1
    no shutdown
  no shutdown
config > gmpls-tunnel-group 3
  type head-end
  far-end remote-uni-c-router-id
  mode protection
  member 1 create
    glsp session-name lsp-name:path-name1
    no shutdown
  no shutdown
config > gmpls-tunnel-group 4
  type head-end
  far-end remote-uni-c-router-id
  mode protection
  member 1 create
    glsp session-name lsp-name:path-name1
    no shutdown
```

```
no shutdown
```

For the tail end node:

```
config > gmpls-tunnel-group 2
  type tail-end
  far-end remote-uni-c-router-id
  mode protection
  member 1 create
    glsp session-name lsp-name:path-name1
    no shutdown
  no shutdown
config > gmpls-tunnel-group 3
  type tail-end
  far-end remote-uni-c-router-id
  mode protection
  member 1 create
    glsp session-name lsp-name:path-name1
    no shutdown
  no shutdown
config > gmpls-tunnel-group 4
  type tail-end
  far-end remote-uni-c-router-id
  mode protection
  member 1 create
    glsp session-name lsp-name:path-name1
    no shutdown
  no shutdown
```

A shutdown of a working path does not trigger a switchover to the protect path. The user should either use the `tools>perform>router>gmpls force` or manual commands, or shutdown the TE-Link, data bearer, or port associated with the gLSP path.

3.16 GMPLS Tunnel Groups

A GMPLS tunnel group is a bundle of gLSPs providing an abstraction of the data bearers that are intended to be associated to one IP interface. This object allows, for example, end-to-end load balancing across the set of data bearers corresponding to a set of gLSPs. A gLSP is bound to a GMPLS tunnel group by a gLSP tunnel (session) name at both the head end and the tail end UNI-C nodes of the gLSP. A sender address (the far-end) may optionally be configured for the tail end of a gLSP in case different head end nodes use overlapping gLSP tunnel names.

```
config
  gmpls-tun-grp gmpls-tun-grp-id
    type {head-end | tail-end}
    far-end remote-uni-c-router-id
    mode {load-sharing | active-standby}
    no mode
    [no] member-threshold threshold [action down]
```

```

member mem-id [create]
    glsp session-name name
    no glsp session-name name
    [no] shutdown
...
[no] shutdown

```

gmpls-tun-grp-id is an unsigned integer from 1 to 1024, shared with the Ethernet tunnel ID range.

The GMPLS Tunnel Group must be configured as either at both the **head-end** or **tail-end** of a set of member gLSPs (identified using the **head-end** or **tail-end** keywords). These keywords are mutually exclusive.

Nodes at the head-end initiate signaling of gLSPs. The **far-end** is the far end of the GMPLS tunnel group. If this node is a head end, then the far end address is taken as the to address for the member gLSPs. Each gLSP that is bound to the tunnel group must have a to address matching the far end address. A binding is held down if a gLSP to and the tunnel group to do not match.

Nodes at the tail end wait for the first path message for a gLSP. The **far-end-address** address must be configured at the tail end. It is the GMPLS Router ID of the head-end UNI-C (the *remote-uni-c-node-id*), and must be configured at the tail end UNI-C of a gLSP. The combination of *session-name* and *remote-uni-c-node-id* provides a unique key to bind an incoming gLSP setup request to a tunnel group. A binding to the tunnel group is held down at the tail end until a gLSP PATH message with a matching *session-name* and source address that matches the tunnel group's far-end address is received.

At the tail end, the **session-name** is composed of the LSP name and Path name as configured at the head end

If **load-sharing** is configured, then all of the gLSPs must terminate on the same far-end node. All of the ports used by gLSPs in a load-sharing must be equivalent in that they must have the same named QoS policy, bandwidth, and so on. Once more than one gLSP is associated with a tunnel group, the QoS policy/scheduler policy cannot be changed in any of the ports. All gLSPs must be unprotected end-to-end in load-sharing mode. Segment protection is allowed for gLSPs associated in load sharing mode to a GMPLS tunnel group.

In **active-standby** mode, only one member gLSP can be associated with the tunnel group.

All members of a tunnel group must be of the same bandwidth.

The **member-threshold** is the number of member gLSPs that must be operationally up before the gmpls tunnel group is considered operationally up.

A member of a GMPLS tunnel group may be treated as down for one of the following reasons. These reason codes are recorded in the `tmnxGmplsTunGrpMemberTable` in the MIB:

- `adminDn` — The member or the related tunnel-grp is administratively down.
- `wpLspDn` — The associated working Lsp-path is down.
- `wpPortDn` — The data-bearer port associated with the working Lsp-path is down.
- `wpPortNoRsrc` — The data-bearer port associated with the working Lsp-path has no resource to support the services over the gmpls-tunnel-grp logical port.
- `ppLspDn` — The associated protect Lsp-path is down.
- `ppPortDn` — The data-bearer port associated with the protect Lsp-path is down.
- `ppPortNoRsrc` — The data-bearer port associated with the protect Lsp-path has no resource to support the services over the gmpls-tunnel-grp logical port.

Note that in the case of `wpPortNoRsrc` and `ppPortNoRsrc`, the term 'resources' relates to QoS or ACL related resources. For example, this can happen when a subsequent physical or data bearing port is added to a GMPLS tunnel group, which already has services running over it. If the new-complex doesn't have the resources to support those services over that GMPLS tunnel group, the related member `operState` would be down with `reasonCode PortNoRsrc`. If a gLSP is already established on a data bearer when a resource failure is experienced, the RSVP PATH message A-Bit is updated so that both ends ensure that the LSP Path is held down.

The user should free resources from the complex, and shutdown/no shutdown the GMPLS tunnel group member. This repeats the resource check, which will bring the member `operUp` if it passes.

A gLSP associated with a tunnel group member will be down if the member is operationally down, or a fault is detected on the associated data bearer.

If a member is in the admin down state, a gLSP will not be set-up. If a gLSP is already up, the RSVP Path message A-Bit is updated so that both ends of the gLSP path are kept down.

3.17 Configuring IP and MPLS in an Overlay Network to Use a GMPLS LSP

IP and MPLS is able to use GMPLS LSPs as transport by bringing a numbered or unnumbered IP interface to an endpoint of one or more gLSPs. This IP interface appears as any other IP interface bound to a network port. The IP interface is bound to the GMPLS tunnel group by a GMPLS tunnel group number configured in the **port** command.

The GMPLS tunnel group number must correspond to a locally configured GMPLS tunnel group.

The following CLI tree illustrates where the GMPLS tunnel group is referenced. This must be done at the nodes at the tunnel groups at both ends of the transport service.

```
config
  router
    interface if-name
      address a.b.c.d | ipv6-address
      port gmpls-tunnel-group gmpls-tunnel-group-id
```

3.18 Configuration Notes

This section describes GMPLS caveats.

- Interfaces must already be configured in the config>router>interface context before they can be specified in GMPLS.
- A router interface must be specified in the config>router>mpls context in order 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.

3.19 GMPLS Configuration Command Reference

3.19.1 Command Hierarchies

- [LMP Commands](#)
- [GMPLS Commands](#)
- [GMPLS Tunnel Group Commands](#)

3.19.1.1 LMP Commands

```

config
  — router
    — [no] lmp
      — gmpls-loopback-address ip-address
      — no gmpls-loopback-address
      — [no] peer peer-node-id
        — [no] control-channel lmp-cc-id
          — hello [interval hello-interval] dead-interval hello-dead-interval
          — hello interval hello-interval [dead-interval hello-dead-interval]
          — peer-interface-address ip-address
          — setup-role {active | passive}
          — ttl limit
          — no ttl
          — [no] shutdown
        — peer-loopback-address ip-address
        — no peer-loopback-address
        — retransmission-interval milliseconds
        — retry-limit limit
        — no retry-limit
        — [no] te-link te-link-id
        — [no] shutdown
      — [no] te-link te-link-id
        — [no] data-bearer data-bearer-id
          — [no] port port-id
          — remote-id remote-id
          — no remote-id
          — [no] shutdown
        — link-name te-link-name
        — no link-name
        — remote-id remote-id
        — no remote-id
        — [no] shutdown
      — [no] shutdown

```

3.19.1.2 GMPLS Commands

```

config
— router
— [no] gmpls
— gr-helper-time max-recovery recovery-interval seconds max restart
  restart-interval seconds
— no gr-helper-time
— keep-multiplier number
— no keep-multiplier
— [no] lsp lsp-name
— e2e-protection-type protection-type
— no e2e-protection-type
— encoding-type encoding-type
— no encoding-type
— generalized-pid generalized-pid
— no generalized-pid
— protect-path path-name
— no protect-path
— bandwidth signal-type signal-type
— no bandwidth
— exclude-srlg group-name [group-name ... (up to 5 max)]
— no exclude-srlg [group-name [group-name ... (up to 5 max)]]
— peer-node peer-node-id
— no peer-node
— segment-protection-type protection-type
— no segment-protection-type
— [no] shutdown
— retry-limit retry-limit
— no retry-limit
— retry-timer seconds
— no retry-timer
— revert-timer seconds
— no revert-timer
— [no] shutdown
— switching-type switching-type
— no switching-type
— to ip-address
— working-path path-name
— no working-path
— bandwidth signal-type signal-type
— no bandwidth
— exclude-srlg group-name [group-name ... (up to 5 max)]
— no exclude-srlg [group-name [group-name ... (up to 5 max)]]
— peer-node peer-node-id
— no peer-node
— segment-protection-type protection-type
— no segment-protection-type
— [no] shutdown
— lsp-init-retry-timeout seconds
— no lsp-init-retry-timeout
— path path-name
— no path

```

- [no] **shutdown**
- **hop** *hop-index* *node-id* *node-id* [*te-link* *te-link-id*] {*strict* | *loose*}
- no **hop** *hop-index*
- [no] **peer** *peer-node-id*
 - [no] **hello-interval** *hello-interval*
 - [no] **lsp-hold-timer** *hold-time*
 - [no] **shutdown**
- **rapid-retransmit-time** *hundred-milliseconds*
- no **rapid-retransmit-time**
- **rapid-retry-limit** *limit*
- no **rapid-retry-limit**
- **refresh-time** *seconds*
- no **refresh-time**
- [no] **shutdown**
 - [no] **te-link** *te-link-id*
 - [no] **shutdown**

3.19.1.3 GMPLS Tunnel Group Commands

- ```
config
— [no] gmpls-tun-grp gmpls-tunnel-group-id
 — description description-string
 — no description
 — far-end ip-address
 — no far-end
 — [no] member member-id
 — [no] glsp session-name name
 — [no] shutdown
 — [no] member-threshold threshold
 — mode mode
 — [no] shutdown
 — type [head-end | tail-end]
```

## 3.19.2 Command Descriptions

### 3.19.2.1 LMP Commands

Imp

**Syntax** [no] Imp

**Context** config>router

---

|                    |                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command creates a context for the configuration of the Link Management Protocol (LMP) on the system. |
|--------------------|-----------------------------------------------------------------------------------------------------------|

|                |        |
|----------------|--------|
| <b>Default</b> | no lmp |
|----------------|--------|

## gmpls-loopback-address

|               |                                                                                     |
|---------------|-------------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>gmpls-loopback-address</b> <i>ip-address</i><br><b>no gmpls-loopback-address</b> |
|---------------|-------------------------------------------------------------------------------------|

|                |                   |
|----------------|-------------------|
| <b>Context</b> | config>router>lmp |
|----------------|-------------------|

|                    |                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command specifies the GMPLS Loopback Address to be used by LMP. A corresponding <b>gmpls-loopback</b> interface must have been configured for LMP to be enabled. |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|                |                           |
|----------------|---------------------------|
| <b>Default</b> | no gmpls-loopback-address |
|----------------|---------------------------|

|                   |                                                |
|-------------------|------------------------------------------------|
| <b>Parameters</b> | <i>ip-address</i> — Specifies an IPv4 address. |
|-------------------|------------------------------------------------|

## peer

|               |                                      |
|---------------|--------------------------------------|
| <b>Syntax</b> | <b>[no] peer</b> <i>peer-node-id</i> |
|---------------|--------------------------------------|

|                |                   |
|----------------|-------------------|
| <b>Context</b> | config>router>lmp |
|----------------|-------------------|

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command creates a context to enable the specification of the LMP peer parameters. It also specifies the LMP peer node. For a GMPLS UNI, this is the UNI-N node at the far end of the IP control channel for the GMPLS UNI. If the peer loopback address is entered using the <b>peer-loopback-address</b> command, then this is used as the routable peer address; otherwise the <i>peer-node-id</i> is assumed to correspond to a routable peer loopback |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|                |         |
|----------------|---------|
| <b>Default</b> | no peer |
|----------------|---------|

|                   |                                                                                                                                |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>peer-node-id</i> — An identifier for the LMP peer node. This may be an IPv4-formatted address or a 32-bit unsigned integer. |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------|

|               |                           |
|---------------|---------------------------|
| <b>Values</b> | a.b.c.d   1 to 4294967295 |
|---------------|---------------------------|

## control-channel

|               |                                              |
|---------------|----------------------------------------------|
| <b>Syntax</b> | <b>[no] control-channel</b> <i>lmp-cc-id</i> |
|---------------|----------------------------------------------|

|                |                        |
|----------------|------------------------|
| <b>Context</b> | config>router>lmp>peer |
|----------------|------------------------|

|                    |                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command enables the context for configuring an IP control channel for use by GMPLS UNI control plane (RSVP and LMP). |
|--------------------|---------------------------------------------------------------------------------------------------------------------------|

|                |                    |
|----------------|--------------------|
| <b>Default</b> | no control-channel |
|----------------|--------------------|

---

**Parameters** *lmp-cc-id* — An unsigned integer identifier for the control channel.

**Values** 1 to 42949672

## hello

**Syntax** **hello** [*interval hello-interval*] *dead-interval* *hello-dead-interval*  
**hello** *interval hello-interval* [*dead-interval hello-dead-interval*]

**Context** config>router>lmp>peer>control-channel

**Description** This command configures the transmission interval for LMP Hello packets. The **dead-interval** specifies the period after which the IPCC is declared down if no hello packets are received from the LMP peer.

**Default** n/a

**Parameters** *interval hello-interval* — The interval at which LMP hello packets are sent on an IP control channel.

**Values** 1000 to 65535 ms

**Default** 1000 ms

*dead-interval hello-dead-interval* — The interval after which the IPCC is declared down if no hello packets are received from the LMP peer.

**Values** 1000 to 65535 ms

**Default** 1000 ms

## peer-interface-address

**Syntax** **peer-interface-address** *ip-address*

**Context** config>router>lmp>peer>control-channel

**Description** This command configures the mandatory **peer-interface-address**. It is the destination address of the IPCC on the peer UNI-N used to reach the GMPLS Router ID of the UNI-N peer. It corresponds to the *lmpCcRemotepAddr* in RFC 4631.

**Default** n/a

**Parameters** *ip-address* — The interface address of the IPCC next-hop.

**Values** *ipv4-address* — a.b.c.d  
*ipv6-address* — x:x:x:x:x:x:x (eight 16-bit pieces)  
x:x:x:x:x:d.d.d.d  
x to [0 to FFFF]H  
d to [0 to 255]D

---

## setup-role

|                    |                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>setup-role {active   passive}</b>                                                                                                        |
| <b>Context</b>     | config>router>lmp>peer>control-channel                                                                                                      |
| <b>Description</b> | This command specifies whether this node takes the active or the passive role in establishing the LMP session to the peer over a GMPLS UNI. |
| <b>Default</b>     | n/a                                                                                                                                         |
| <b>Parameters</b>  | <b>active</b> — The router takes the active role. (Default)<br><b>passive</b> — The router takes the passive role.                          |

## ttl

|                    |                                                                                                                                                                                                                                |               |          |                |   |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|----------|----------------|---|
| <b>Syntax</b>      | <b>ttl <i>limit</i></b><br><b>no ttl</b>                                                                                                                                                                                       |               |          |                |   |
| <b>Context</b>     | config>router>lmp>peer>control-channel                                                                                                                                                                                         |               |          |                |   |
| <b>Description</b> | This command configures the time to live (TTL) for all packets (GMPLS RSVP and LMP) on the IP control channel, which allows the TTL to be optimized for multiple-hop communication networks between the GMPLS UNI-C and UNI-N. |               |          |                |   |
| <b>Default</b>     | no ttl                                                                                                                                                                                                                         |               |          |                |   |
| <b>Parameters</b>  | <b>limit</b> — The TTL value for the packets.<br><table><tr><td><b>Values</b></td><td>2 to 255</td></tr><tr><td><b>Default</b></td><td>1</td></tr></table>                                                                     | <b>Values</b> | 2 to 255 | <b>Default</b> | 1 |
| <b>Values</b>      | 2 to 255                                                                                                                                                                                                                       |               |          |                |   |
| <b>Default</b>     | 1                                                                                                                                                                                                                              |               |          |                |   |

## shutdown

|                    |                                                                           |
|--------------------|---------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                                      |
| <b>Context</b>     | config>router>lmp>peer>control-channel                                    |
| <b>Description</b> | This command administratively enables or disables the IP control channel. |
| <b>Default</b>     | no shutdown                                                               |

## peer-loopback-address

|               |                                                                                   |
|---------------|-----------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>peer-loopback-address <i>ip-address</i></b><br><b>no peer-loopback-address</b> |
|---------------|-----------------------------------------------------------------------------------|

|                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>                                                                    | config>router>Imp>peer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>                                                                | <p>The IP address corresponding to the GMPLS loopback address configured on the LMP peer. If <b>peer-loopback-address</b> is entered, then this is used as the routable peer address, otherwise the <i>peer-node-id</i> is assumed to correspond to a routable peer loopback.</p> <p><b>peer-loopback-address</b> is an optional configurable field. If <b>peer-loopback-address</b> is not configured, the router will use <b>Imp-peer-node-id</b> (i.e. LmpNbrNodeId as per RFC 4631) as the dstIpAddr in the IP-header for the peer-specific messages (that is, Link summary msgs, RSVP msgs).</p> |
|  | <p><b>Note:</b> The <b>peer-interface-address</b> is mandatory; it is the destination address of the IPCC on the peer UNI-N used to reach the GMPLS Router ID of the UNI-N peer. It corresponds to the ImpCcRemoteIpAddr in RFC 4631.</p>                                                                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>                                                                    | no peer-loopback-address                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>                                                                 | <p><i>ip-address</i> — The GMPLS control plane loopback address of the IPCC next-hop.</p> <p><b>Values</b></p> <ul style="list-style-type: none"><li>ipv4-address — a.b.c.d</li><li>ipv6-address — x:x:x:x:x:x:x (eight 16-bit pieces)</li><li>x:x:x:x:x:d.d.d.d</li><li>x to [0 to FFFF]H</li><li>d to [0 to 255]D</li></ul>                                                                                                                                                                                                                                                                         |

## retransmission-interval

|                    |                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>retransmission-interval</b> <i>milliseconds</i>                                                                                        |
| <b>Context</b>     | config>router>Imp>peer                                                                                                                    |
| <b>Description</b> | This command specifies the interval between resubmitted LMP messages.                                                                     |
| <b>Default</b>     | n/a                                                                                                                                       |
| <b>Parameters</b>  | <p><i>milliseconds</i> — Specifies the retransmission interval, in ms.</p> <p><b>Values</b> 1 to 4294967295</p> <p><b>Default</b> 500</p> |

## retry-limit

|                |                                                          |
|----------------|----------------------------------------------------------|
| <b>Syntax</b>  | <b>retry-limit</b> <i>limit</i><br><b>no retry-limit</b> |
| <b>Context</b> | config>router>Imp>peer                                   |

---

|                    |                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------|
| <b>Description</b> | This command specifies how many times LMP resends a message before restarting the process. |
| <b>Default</b>     | no retry-limit                                                                             |
| <b>Parameters</b>  | <i>limit</i> — Specifies the number of reattempts.                                         |
| <b>Values</b>      | 1 to 4294967295                                                                            |

te-link

|                    |                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] te-link <i>te-link-id</i></b>                                                                                                                                                            |
| <b>Context</b>     | config>router>Imp>peer                                                                                                                                                                           |
| <b>Description</b> | This command assigns a Traffic Engineering (TE) Link to a given LMP peer. The TE Link with ID <i>te-link-id</i> must already have been created under <b>config&gt;router&gt;Imp&gt;te-link</b> . |
| <b>Default</b>     | no te-link                                                                                                                                                                                       |
| <b>Parameters</b>  | <i>te-link-id</i> — Specifies the ID of the TE Link.                                                                                                                                             |
| <b>Values</b>      | 1 to 4294967295   <i>te-link-name</i><br><i>te-link-name</i> : 32 character (max) name of the TE Link                                                                                            |

shutdown

|                    |                                                                          |
|--------------------|--------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                                     |
| <b>Context</b>     | config>router>Imp>peer                                                   |
| <b>Description</b> | This command administratively enables or disables LMP with a given peer. |
| <b>Default</b>     | no shutdown                                                              |

te-link

|                    |                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] te-link <i>te-link-id</i></b>                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>Imp                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | This command creates a Traffic Engineering (TE) Link in LMP across a GMPLS UNI. An unsigned integer TE link ID must be specified when the TE Link is first created. Once the link is created, the user can configure the link name (i.e. 'link-name te-link-name'). From here, the user can refer to this TE Link by either the unsigned integer or the ASCII name. |
| <b>Default</b>     | no te-link                                                                                                                                                                                                                                                                                                                                                          |



---

|                   |                                                                                                       |
|-------------------|-------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>te-link-id</i> — Specifies the ID of the TE Link.                                                  |
| <b>Values</b>     | 1 to 4294967295   <i>te-link-name</i><br><i>te-link-name</i> : 32 character (max) name of the TE Link |

data-bearer

|                    |                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] data-bearer <i>data-bearer-id</i></b>                                                                            |
| <b>Context</b>     | config>router>Imp>te-link                                                                                                |
| <b>Description</b> | This command creates a data bearer assigned to a TE Link. Only one data bearer may be configured within a given TE Link. |
| <b>Default</b>     | no data-bearer                                                                                                           |
| <b>Parameters</b>  | <i>data-bearer-id</i> — Specifies the ID of the data bearer.                                                             |
| <b>Values</b>      | 1 to 4294967295                                                                                                          |

port

|                    |                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] port <i>port-id</i></b>                                                                                              |
| <b>Context</b>     | config>router>Imp>te-link>data-bearer                                                                                        |
| <b>Description</b> | This command configures the port associated with the data bearer. The port must be a physical black and white Ethernet port. |
| <b>Default</b>     | no port                                                                                                                      |
| <b>Parameters</b>  | <i>port-id</i> — Specifies the ID of the port.                                                                               |
| <b>Values</b>      |                                                                                                                              |
| <i>port-id</i>     | <i>slot/mda/port</i> [.channel]                                                                                              |
| eth-sat-id         | esat-id/slot/port                                                                                                            |
|                    | esat                                                                                                                         |
|                    | id                                                                                                                           |
|                    | keyword                                                                                                                      |
|                    | 1 to 20                                                                                                                      |
| pxc-id             | pxc-id.sub-port                                                                                                              |
|                    | pxc                                                                                                                          |
|                    | keyword                                                                                                                      |
|                    | id                                                                                                                           |
|                    | 1 to 64                                                                                                                      |
|                    | sub-port                                                                                                                     |
|                    | a, b                                                                                                                         |

remote-id

|               |                                   |
|---------------|-----------------------------------|
| <b>Syntax</b> | <b>remote-id <i>remote-id</i></b> |
|---------------|-----------------------------------|

**no remote-id**

|                    |                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>Imp>te-link>data-bearer                                                                                             |
| <b>Description</b> | This command configures the identifier assigned to the data-bearer at the LMP peer node. For a GMPLS UNI, this is the UNI-N node. |
| <b>Default</b>     | no remote-id                                                                                                                      |
| <b>Parameters</b>  | <i>remote-id</i> — Specifies the ID of the data-bearer at the LMP peer node.                                                      |

**Values** 1 to 4294967295

## shutdown

|                    |                                                                    |
|--------------------|--------------------------------------------------------------------|
| <b>Syntax</b>      | [no] shutdown                                                      |
| <b>Context</b>     | config>router>Imp>te-link>data-bearer                              |
| <b>Description</b> | This command administratively enables or disables the data bearer. |
| <b>Default</b>     | no shutdown                                                        |

## link-name

|                    |                                                                |
|--------------------|----------------------------------------------------------------|
| <b>Syntax</b>      | <b>link-name</b> <i>te-link-name</i><br><b>no link-name</b>    |
| <b>Context</b>     | config>router>Imp>te-link                                      |
| <b>Description</b> | This command configures text names for the TE Link.            |
| <b>Default</b>     | n/a                                                            |
| <b>Parameters</b>  | <i>te-link-name</i> — Specifies the text name for the TE Link. |

**Values** 32 characters maximum text string

## remote-id

|                    |                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>remote-id</b> <i>id</i><br><b>no remote-id</b>                                                                             |
| <b>Context</b>     | config>router>Imp>te-link                                                                                                     |
| <b>Description</b> | This command configures the identifier assigned to the TE Link at the LMP peer node. For a GMPLS UNI, this is the UNI-N node. |
| <b>Default</b>     | no remote-id                                                                                                                  |

---

|                   |                                                                     |
|-------------------|---------------------------------------------------------------------|
| <b>Parameters</b> | <i>id</i> — Specifies the identifier for the LMP peer node TE Link. |
| <b>Values</b>     | 1 to 4294967295                                                     |

## shutdown

|                    |                                                                |
|--------------------|----------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                           |
| <b>Context</b>     | config>router>Imp>te-link                                      |
| <b>Description</b> | This command administratively enables or disables the TE Link. |
| <b>Default</b>     | no shutdown                                                    |

## shutdown

|                    |                                                        |
|--------------------|--------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                   |
| <b>Context</b>     | config>router>Imp                                      |
| <b>Description</b> | This command administratively enables or disables LMP. |
| <b>Default</b>     | no shutdown                                            |

### 3.19.2.2 GMPLS Commands

## gmpls

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] gmpls</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>router                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command enables the context to configure GMPLS parameters. GMPLS is not enabled by default and must be explicitly enabled using <b>no shutdown</b>. The <b>shutdown</b> command administratively disables GMPLS.</p> <p>The <b>no</b> form of this command deletes this GMPLS protocol instance; this will remove all configuration parameters for this GMPLS instance.</p> <p>GMPLS must be shut down before the GMPLS instance can be deleted. If GMPLS is not shut down when the <b>no gmpls</b> command is executed, a warning message on the console indicates that GMPLS is still administratively up.</p> |
| <b>Default</b>     | no gmpls                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

---

gr-helper-time

|                    |                                                                                                                                                                                                                                                                                                                                                                                                      |               |           |                |     |               |          |                |     |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|-----------|----------------|-----|---------------|----------|----------------|-----|
| <b>Syntax</b>      | <b>gr-helper-time max-recovery</b> <i>recovery-interval</i> <b>seconds max restart</b> <i>restart-interval</i> <b>seconds</b><br><b>no gr-helper-time</b>                                                                                                                                                                                                                                            |               |           |                |     |               |          |                |     |
| <b>Context</b>     | config>router>gmpls                                                                                                                                                                                                                                                                                                                                                                                  |               |           |                |     |               |          |                |     |
| <b>Description</b> | <p>This command configures the local values for the max-recovery and the max-restart intervals used in the RSVP Graceful Restart Helper feature when applied to a GMPLS UNI.</p> <p>The values are configured globally in GMPLS.</p> <p>The <b>no</b> version of this command re-instates the default value for the delay timer.</p>                                                                 |               |           |                |     |               |          |                |     |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                                                                                                                                                                  |               |           |                |     |               |          |                |     |
| <b>Parameters</b>  | <p><i>recovery-interval</i> — Specifies the maximum recovery interval value, in s.</p> <table><tr><td><b>Values</b></td><td>1 to 1800</td></tr><tr><td><b>Default</b></td><td>300</td></tr></table> <p><i>restart-interval</i> — Specifies the maximum restart interval value, in s.</p> <table><tr><td><b>Values</b></td><td>1 to 300</td></tr><tr><td><b>Default</b></td><td>180</td></tr></table> | <b>Values</b> | 1 to 1800 | <b>Default</b> | 300 | <b>Values</b> | 1 to 300 | <b>Default</b> | 180 |
| <b>Values</b>      | 1 to 1800                                                                                                                                                                                                                                                                                                                                                                                            |               |           |                |     |               |          |                |     |
| <b>Default</b>     | 300                                                                                                                                                                                                                                                                                                                                                                                                  |               |           |                |     |               |          |                |     |
| <b>Values</b>      | 1 to 300                                                                                                                                                                                                                                                                                                                                                                                             |               |           |                |     |               |          |                |     |
| <b>Default</b>     | 180                                                                                                                                                                                                                                                                                                                                                                                                  |               |           |                |     |               |          |                |     |

## keep-multiplier

|                    |                                                                                                                                                                                                |               |          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|----------|
| <b>Syntax</b>      | <b>keep multiplier</b> <i>number</i><br><b>no keep-multiplier</b>                                                                                                                              |               |          |
| <b>Context</b>     | config>router>gmpls                                                                                                                                                                            |               |          |
| <b>Description</b> | <p>This command configures the integer used by RSVP to declare that a reservation is down or the neighbor is down.</p> <p>The <b>no</b> form of this command reverts to the default value.</p> |               |          |
| <b>Default</b>     | 3                                                                                                                                                                                              |               |          |
| <b>Parameters</b>  | <p><i>number</i> — Specifies the keep multiplier value.</p> <table><tr><td><b>Values</b></td><td>1 to 255</td></tr></table>                                                                    | <b>Values</b> | 1 to 255 |
| <b>Values</b>      | 1 to 255                                                                                                                                                                                       |               |          |

## lsp

|               |                                 |
|---------------|---------------------------------|
| <b>Syntax</b> | <b>[no] lsp</b> <i>lsp-name</i> |
|---------------|---------------------------------|

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>gmpls                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command creates a GMPLS LSP that is signaled dynamically by the router.</p> <p>When the LSP is created, the egress router must be specified using the <b>to</b> command and a <b>working-path</b> must be specified.</p> <p>GMPLS LSPs are created in the administratively down (<b>shutdown</b>) state.</p> <p>The <b>no</b> form of this command deletes the GMPLS LSP. All configuration information associated with this GMPLS LSP is lost. The GMPLS LSP must be administratively shut down before it can be deleted.</p> |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b>  | <i>lsp-name</i> — Specifies the identifier for the GMPLS LSP. The LSP name can be up to 32 characters long and must be unique.                                                                                                                                                                                                                                                                                                                                                                                                         |

## e2e-protection-type

|                    |                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>e2e-protection-type <i>protection-type</i></b><br><b>no e2e-protection-type</b>                                                                                                                                                                                                              |
| <b>Context</b>     | config>router>gmpls>lsp                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command defines the end-to-end recovery type for the GLSP. This is the recovery model between the source and terminating UNI-C nodes of the GMPLS LSP.</p> <p>The <b>no</b> form of this command removes any configured end-to-end recovery, and the GMPLS LSP becomes unprotected.</p> |
| <b>Default</b>     | no e2e-protection-type                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <i>protection-type</i> — Specifies the end-to-end GMPLS recovery type.                                                                                                                                                                                                                          |
| <b>Values</b>      | {unprotected   1toN   sbr}                                                                                                                                                                                                                                                                      |
| <b>Default</b>     | unprotected                                                                                                                                                                                                                                                                                     |

## encoding-type

|                    |                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>encoding-type <i>encoding-type</i></b><br><b>no encoding-type</b>                                                       |
| <b>Context</b>     | config>router>gmpls>lsp                                                                                                    |
| <b>Description</b> | This command configures the encoding type of the payload carried by the GMPLS LSP. <b>line</b> is the only supported type. |
| <b>Default</b>     | no encoding-type                                                                                                           |

---

|                   |                                                     |
|-------------------|-----------------------------------------------------|
| <b>Parameters</b> | <i>encoding-type</i> — Specifies the encoding type. |
| <b>Values</b>     | line                                                |
| <b>Default</b>    | line                                                |

## generalized-pid

|                    |                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>generalized-pid <i>generalized-pid</i></b><br><b>no generalized-pid</b>                                                                                                                              |
| <b>Context</b>     | config>router>gmpls>lsp                                                                                                                                                                                 |
| <b>Description</b> | This command configures the type of payload carried by the gLSP. Standard ethertype values are used for packet and Ethernet LSPs (see RFC 3471). Only Ethernet (value 33) is supported in Release 13.0. |
| <b>Default</b>     | <b>no generalized-pid</b>                                                                                                                                                                               |
| <b>Parameters</b>  | <i>generalized-pid</i> — Specifies the name of the generalized-pid.                                                                                                                                     |
| <b>Values</b>      | ethernet                                                                                                                                                                                                |
| <b>Default</b>     | ethernet                                                                                                                                                                                                |

## retry-limit

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>retry-limit <i>retry-limit</i></b><br><b>no retry-limit</b>                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>router>gmpls>lsp                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This optional command specifies the number of attempts software should make to re-establish the GMPLS LSP after it has failed. After each successful attempt, the counter is reset to zero.</p> <p>When the specified number is reached, no more attempts are made and the GMPLS LSP path is put into the <b>shutdown</b> state.</p> <p>Use the <b>config router gmpls lsp <i>lsp-name</i> no shutdown</b> command to bring up the path after the retry limit is exceeded.</p> |
| <b>Default</b>     | 0 (no limit, retries forever)                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <i>retry-limit</i> — Specifies the number of retries.                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Values</b>      | 0 to 10000                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Default</b>     | 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## retry-timer

|                    |                                                                                                                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>retry-timer</b> <i>seconds</i><br><b>no retry-timer</b>                                                                                                                                                                                         |
| <b>Context</b>     | config>router>gmpls>lsp                                                                                                                                                                                                                            |
| <b>Description</b> | This command configures the time ( in s), for LSP re-establishment attempts after it has failed. The retry time is jittered to +/- 25% of its nominal value.<br><br>The <b>no</b> form of this command reverts the parameter to the default value. |
| <b>Default</b>     | 30                                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <i>seconds</i> — Specifies the amount of time (in s), between attempts to re-establish the LSP after it has failed.<br><br><b>Values</b> 0 to 600 s<br><b>Default</b> 30                                                                           |

## revert-timer

|                    |                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>revert-timer</b> <i>seconds</i><br><b>no revert-timer</b>                                                                                                              |
| <b>Context</b>     | config>router>gmpls>lsp                                                                                                                                                   |
| <b>Description</b> | This command configures the time (in s), for LSP reversion attempts after it has failed.<br><br>The <b>no</b> form of the command reverts the timer to the default value. |
| <b>Default</b>     | 0                                                                                                                                                                         |
| <b>Parameters</b>  | <i>seconds</i> — Specifies the time (in s), for the LSP to attempt reversion after failure.<br><br><b>Values</b> 0 to 1800<br><b>Default</b> 0                            |

## shutdown

|                    |                                                                  |
|--------------------|------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                             |
| <b>Context</b>     | config>router>gmpls>lsp                                          |
| <b>Description</b> | This command administratively enables or disables the GMPLS LSP. |
| <b>Default</b>     | shutdown                                                         |

---

switching-type

|                    |                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>switching-type</b> <i>switching-type</i><br><b>no switching-type</b>                                                                                                                                                                                              |
| <b>Context</b>     | config>router>gmpls>lsp                                                                                                                                                                                                                                              |
| <b>Description</b> | This command configures the type of switching required for the gLSP. As defined in RFC 3471. The default CLI value is ethernet, which indicates that Digital Channel Switch Capable (DCSC) should be signaled. Ethernet is the only supported value in Release 13.0. |
| <b>Default</b>     | ethernet                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <i>switching-type</i> — Specifies the required type of switching.<br><b>Values</b> ethernet                                                                                                                                                                          |

## to

|                    |                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>to</b> <i>ip-address</i>                                                                                                                          |
| <b>Context</b>     | config>router>gmpls>lsp                                                                                                                              |
| <b>Description</b> | This command specifies the GMPLS loopback address of the far-end UNI-C router for a GMPLS LSP. When creating a GMPLS LSP, this command is mandatory. |
| <b>Parameters</b>  | <i>ip-address</i> — Specifies the system IP address of the far-end UNI-C router.                                                                     |

## working-path

|                    |                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>working-path</b> <i>path-name</i><br><b>no working-path</b>                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>router>gmpls>lsp                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | This command specifies the working path for a GMPLS LSP. One working path must be specified for each GMPLS LSP. The path-name parameter must correspond to a path defined under <b>config&gt;router&gt;gmpls&gt;path</b> .<br><br>The <b>no</b> form of the command removes the working-path definition. |
| <b>Default</b>     | no working-path                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <i>path-name</i> — Specifies the name of the path used by the working path.<br><b>Values</b> 32 characters maximum text string                                                                                                                                                                           |



---

## protect-path

|                    |                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>protect-path</b> <i>path-name</i><br><b>no protect-path</b>                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>router>gmpls>lsp                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command specifies the protect path for a GMPLS LSP. At least one protect path must be specified if a GMPLS LSP uses 1-to-N end-to-end protection. The path-name parameter must correspond to a path defined under <b>config&gt;router&gt;gmpls&gt;path</b>.</p> <p>The <b>no</b> form of the command removes the protect-path definition.</p> |
| <b>Default</b>     | no protect-path                                                                                                                                                                                                                                                                                                                                       |
| <b>Parameters</b>  | <i>path-name</i> — The name of the path used by the protect path.<br><b>Values</b> 32 characters maximum text string                                                                                                                                                                                                                                  |

## bandwidth

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>bandwidth</b> <i>signal-type</i> <i>signal-type</i><br><b>no bandwidth</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>router>gmpls>lsp>working-path<br>config>router>gmpls>lsp>protect-path                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command specifies the bandwidth to be signaled for the path of the GMPLS LSP. Bandwidth is specified in terms of the RFC 3471 signal type name.</p> <p>If an empty path is configured or the first hop TE Link is not configured, the system will automatically select a TE Link to use for a GMPLS LSP path based on the lowest available TE Link ID with a matching bandwidth (if a bandwidth is configured for the GMPLS LSP). During a data-bearer link allocation request, an RSVP-requested GMPLS LSP BW can be either a non-zero value as per RFC 3471 signal-type, or it can be zero). There are the following cases:</p> <ul style="list-style-type: none"><li>• Case 1 — The requested BW is non-zero as per RFC 3471 Signal-type config:<ul style="list-style-type: none"><li>a) When a TE (or TE/DB) Link is configured in the related hop, LMP checks whether the related port BW is the same (exact match) as the requested BW, and allocates the port (provided any other checks are OK).</li><li>b) When the related Hop is empty: LMP finds a db-link port to the peer node matching the requested BW, and allocates it.</li></ul></li><li>• Case 2 — Requested BW is zero:<ul style="list-style-type: none"><li>a) When a TE (or TE/DB) Link is configured in the related hop, LMP allocates the port (provided the other checks are OK), and provides the port BW to RSVP to use in signaling.</li></ul></li></ul> |

b) When the related Hop is empty, LMP finds the first available db-link to the peer (based on lower db-link Id), and allocates it and provides the port BW to RSVP to use in signaling.

The **no** form of the command updates the bandwidth to zero.

|                   |                                                                                                                                                                                                                                               |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | 0                                                                                                                                                                                                                                             |
| <b>Parameters</b> | <i>signal-type</i> — Specifies the RFC 3471 name of the signal type representing the requested bandwidth for the GMPLS LSP path.                                                                                                              |
| <b>Values</b>     | {ds0   ds1   e1   ds2   e2   ethernet   e3   ds3   sts-1   fast-ethernet   e4   fc-0-133m   oc-3/stm1   fc-0-266m   fc-0-531m   oc-12/stm-4   gige   fc-0-1062m   oc-48/stm-16   oc-192/stm-64   10gige-ieee   oc-768/stm-256   100gige-ieee} |

## exclude-srlg

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>exclude-srlg</b> <i>group-name</i> [ <i>group-name</i> ... (up to 5 max)]<br><b>no exclude-srlg</b> [ <i>group-name</i> [ <i>group-name</i> ... (up to 5 max)]]                                                                                                                                                                                                                                                   |
| <b>Context</b>     | config>router>gmpls>lsp>working-path<br>config>router>gmpls>lsp>protect-path                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | This command specifies a list of one to five SRLG groups in the optical network which the router can request to the UNI-N that the GMPLS LSP path should avoid by signaling it in the XRO of the RSVP path message. Each <i>group-name</i> must have been defined under <b>config&gt;router&gt;if-attribute&gt;srlg-group</b> .<br><br>The <b>no</b> form of the command removes the list of SRLG groups to exclude. |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <i>group-name</i> — Specifies the name of the SRLG.                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Values</b>      | 32 characters maximum text string                                                                                                                                                                                                                                                                                                                                                                                    |

## peer-node

|                    |                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>peer-node</b> <i>peer-node-id</i><br><b>no peer-node</b>                                                                                                                                                                                                          |
| <b>Context</b>     | config>router>gmpls>lsp>working-path<br>config>router>gmpls>lsp>protect-path                                                                                                                                                                                         |
| <b>Description</b> | This command specifies a peer node to use for the first hop of the GMPLS LSP. If specified, this command forces the GMPLS LSP to use a specific UNI-N node on ingress to the optical network. This command is only applicable if 1toN end to end protection is used. |

The **no** form of the command removes the list of SRLG groups to exclude.

|                   |                                                                                                                          |
|-------------------|--------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | none                                                                                                                     |
| <b>Parameters</b> | <i>peer-node-id</i> — The node ID of the peer UNI-N. This may be an ipv4-formatted address or a 32-bit unsigned integer. |
| <b>Values</b>     | a.b.c.d   1 – 4294967295                                                                                                 |

## segment-protection-type

|                    |                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>segment-protection-type</b> <i>protection-type</i><br><b>no segment-protection-type</b>                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>gmpls>lsp>working-path<br>config>router>gmpls>lsp>protect-path                                                                                                                                                                                                                                                                                |
| <b>Description</b> | This command defines the requested segment recovery type for the GLSP path. This is the recovery capability requested by the router UNI-C to the UNI-N for recovery in segments of the optical network between ingress and egress UNI-N nodes.<br><br>The <b>no</b> form of this command removes the configured segment recovery, reverting to unprotected. |
| <b>Default</b>     | <b>no segment-protection-type</b>                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <i>protection-type</i> — Specifies the requested GMPLS segment recovery type.                                                                                                                                                                                                                                                                               |
| <b>Values</b>      | {unprotected   sbr   gr   sncp   prc}                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>     | unprotected                                                                                                                                                                                                                                                                                                                                                 |

## shutdown

|                    |                                                                              |
|--------------------|------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                                         |
| <b>Context</b>     | config>router>gmpls>lsp>working-path<br>config>router>gmpls>lsp>protect-path |
| <b>Description</b> | This command administratively enables or disables the GMPLS LSP path.        |
| <b>Default</b>     | no shutdown                                                                  |

## lsp-init-retry-timeout

|               |                                                                                  |
|---------------|----------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>lsp-init-retry-timeout</b> <i>seconds</i><br><b>no lsp-init-retry-timeout</b> |
|---------------|----------------------------------------------------------------------------------|

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>gmpls                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command configures the initial GMPLS LSP path retry timer.</p> <p>The new GMPLS LSP path initial retry timer is used instead of the retry-timer to abort the retry cycle when no RESV is received. The retry-timer exclusively governs the time between two retry cycles and to handle retrying of a GMPLS LSP path in a failure case with PATH errors or RESVTear.</p> <p>The <b>no</b> form of this command returns the timer to the default value.</p> |
| <b>Default</b>     | no lsp-init-retry-timeout                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <i>seconds</i> — Specifies the time (in s), between retry cycles.                                                                                                                                                                                                                                                                                                                                                                                                 |

## path

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>path</b> <i>path-name</i><br><b>no path</b> <i>path-name</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>router>gmpls                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command creates the path to be used for a GMPLS LSP. A path can be used by multiple GMPLS LSPs. A path can specify some or all hops from ingress to egress and they can be either strict or loose.</p> <p>Paths are created in a <b>no shutdown</b> 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 GMPLS LSP using the path becomes operationally down.</p> <p>The <b>no</b> form of this command deletes the path and all its associated configuration information. All the GMPLS LSPs that are currently using this path will be affected. A path must be shut down and unbound from all GMPLS LSPs using the path before it can be deleted. The <b>no path</b> <i>path-name</i> command will not result in any action except a warning message on the console indicating that the path may be in use.</p> |
| <b>Parameters</b>  | <i>path-name</i> — Specifies a unique case-sensitive name label for the LSP path.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Values</b>      | 32 characters maximum alphanumeric string                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## shutdown

|                    |                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                                                                                                               |
| <b>Context</b>     | config>router>gmpls>path                                                                                                                           |
| <b>Description</b> | This command disables GMPLS LSPs using the path. All services using these GMPLS LSPs are affected. Paths are created in the <b>shutdown</b> state. |

The **no** form of this command administratively enables the path. All LSPs, where this path is defined as primary or defined as standby secondary, are (re)established.

**Default** no shutdown

## hop

**Syntax** **hop** *hop-index* *node-id* *node-id* [*te-link* *te-link-id*] {**strict** | **loose**}  
**no** **hop** *hop-index*

**Context** config>router>gmpls>path

**Description** This command specifies the node ID of the hops that the GMPLS LSP should traverse on its way to the egress UNI-C router.

The GMPLS LSP ingress and egress node IDs can be included as the first and the last hop. This is necessary when inter-operating with the Nokia 1830 PSS.

The **no** form of this command deletes hop list entries for the path. All of the GMPLS LSPs currently using the path are affected. Additionally, all services actively using these GMPLS LSPs are affected. The path must be shut down first in order to delete the hop from the hop list. The **no hop** *hop-index* command will not result in any action except a warning message on the console indicating that the path is administratively up.

**Default** none

**Parameters** *hop-index* — Specifies the order of the 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

*node-id* — Specified the node ID of the transit GMPLS LSR. This can be an IPv4 address or a 32-bit unsigned integer identifier of the data plane switching node of the adjacent UNI-N.

**loose** — Specifies that the route taken by the GMPLS LSP from the previous hop to this hop can traverse through other LSRs. Multiple hop entries with the *node-id* are flagged as errors. Either the loose or strict keyword must be specified.

**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. If the IP address specified is the interface address, then the LSP must use that interface. 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. Multiple hop entries with the same IP address are flagged as errors. Either the loose or strict keyword must be specified.

---

peer

|                    |                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] peer <i>peer-node-id</i></b>                                                                                                                                                                                                                       |
| <b>Context</b>     | config>router>gmpls                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command specifies parameters for the RSVP session to a neighboring GMPLS UNI-N node. The <i>peer-node-id</i> is the control plane identifier for the adjacent UNI-N node.</p> <p>The <b>no</b> form of this command deletes the configuration.</p> |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><i>peer-node-id</i> — Specifies the control plane node ID of the neighboring GMPLS UNI-N node. This can be an ipv4 address or a 32-bit unsigned integer.</p> <p><b>Values</b> {a.b.c.d   1 to 4294967295}</p>                                           |

## hello-interval

|                    |                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] hello-interval <i>hello-interval</i></b>                                                                                                                                                                                           |
| <b>Context</b>     | config>router>gmpls>peer                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command configures the RSVP hello packet interval (in ms), towards the peer UNI-N node.</p> <p>The <b>no</b> form of this command sets the <b>hello-interval</b> to the default of 3000 ms. A value of 0 disables RSVP hellos.</p> |
| <b>Default</b>     | no hello-interval                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><i>hello-interval</i> — Specifies the RSVP hello packet interval, in ms.</p> <p><b>Values</b> 0 to 6000</p> <p><b>Default</b> 3000</p>                                                                                                  |

## lsp-hold-timer

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] lsp-hold-timer <i>hold-time</i></b>                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>router>gmpls>peer                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>This command specifies the amount of time that the ingress node holds before programming its data plane and declaring a GMPLS LSP up. This occurs anytime the ingress UNI-C node brings up a GMPLS LSP path or reroutes a GMPLS LSP. The <i>hold-time</i> value should be configured to reflect the data path programming time for the optical technology used between the ingress and egress UNI-N nodes.</p> |

---

The **no** form of the command reverts the hold-timer to the default value.

|                   |                                                                 |
|-------------------|-----------------------------------------------------------------|
| <b>Default</b>    | no lsp-hold-timer                                               |
| <b>Parameters</b> | <i>hold-timer</i> — Specifies the ingress node hold time, in s. |
| <b>Values</b>     | 5 to 300                                                        |
| <b>Default</b>    | 60                                                              |

## shutdown

|                    |                                                                                       |
|--------------------|---------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                                                  |
| <b>Context</b>     | config>router>gmpls>peer                                                              |
| <b>Description</b> | This command disables or enables RSVP adjacency with the neighboring UNI-N peer node. |
| <b>Default</b>     | shutdown                                                                              |

## rapid-retransmit-time

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>rapid-retransmit-time <i>hundred-milliseconds</i></b><br><b>no rapid-retransmit-time</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>router>gmpls                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command configures the value of the Rapid Retransmission Interval. It is used in the re-transmission mechanism to handle unacknowledged message_id objects and is based on an exponential back-off timer.</p> <p>Re-transmission interval of a RSVP message with the same message_id = 2 * rapid-retransmit-time interval of time.</p> <p>The node stops re-transmission of unacknowledged RSVP messages:</p> <ul style="list-style-type: none"><li>• if the updated back-off interval exceeds the value of the regular refresh interval, or</li><li>• if the number of re-transmissions reaches the value of the <b>rapid-retry-limit</b> parameter, whichever comes first</li></ul> <p>The Rapid Retransmission Interval must be smaller than the regular refresh interval configured in <b>config&gt;router&gt;gmpls&gt;refresh-time</b>.</p> <p>The <b>no</b> form of this command reverts to the default value.</p> |
| <b>Default</b>     | no rapid-retransmit-time                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

---

|                   |                                                                                                                                                      |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>hundred-milliseconds</i> — Specifies the Rapid Retransmission Interval, in units of 100 ms (for example, enter “6” for a 600 ms retransmit time). |
| <b>Values</b>     | 1 to 100                                                                                                                                             |
| <b>Default</b>    | 5                                                                                                                                                    |

## rapid-retry-limit

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>rapid-retry-limit <i>limit</i></b><br><b>no rapid-retry-limit</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>router>gmpls                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command configures 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 message with the same message_id is retransmitted every 2 * rapid-retransmit-time interval of time. The node will stop retransmission of unacknowledged RSVP messages whenever the updated backoff interval exceeds the value of the regular refresh interval, or the number of retransmissions reaches the value of the <b>rapid-retry-limit</b> parameter, whichever comes first.</p> <p>The <b>no</b> form of this command reverts to the default value.</p> |
| <b>Default</b>     | no rapid-retry-limit                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <i>limit</i> — Specifies the Rapid Retry Limit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Values</b>      | 1 to 6                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Default</b>     | 3                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## refresh-time

|                    |                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>refresh-time <i>seconds</i></b><br><b>no refresh-time</b>                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>router>gmpls                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command configures the interval (in s), between the successive Path and Resv refresh messages. RSVP declares the session down after it misses a consecutive number of refresh messages equal to the configured <b>keep-multiplier</b> number.</p> <p>The <b>no</b> form of this command reverts to the default value.</p> |
| <b>Default</b>     | no refresh-time                                                                                                                                                                                                                                                                                                                   |



---

|                   |                                                                                                    |
|-------------------|----------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>seconds</i> — Specifies the interval (in s), between successive Path and Resv refresh messages. |
| <b>Values</b>     | 1 to 65535                                                                                         |
| <b>Default</b>    | 30                                                                                                 |

## shutdown

|                    |                                         |
|--------------------|-----------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                    |
| <b>Context</b>     | config>router>gmpls                     |
| <b>Description</b> | This command disables or enables GMPLS. |
| <b>Default</b>     | shutdown                                |

## te-link

|                    |                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] te-link <i>te-link-id</i></b>                                                                                                                                                                                                |
| <b>Context</b>     | config>router>gmpls                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command enables the use of a Traffic Engineering (TE) Link (which has previously been configured under <b>config&gt;router&gt;imp</b>) in GMPLS.</p> <p>The <b>no</b> form of this command reverts to the default value.</p> |
| <b>Default</b>     | no te-link                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <p><i>te-link-id</i> — Specifies the ID or name of the configured TE Link.</p> <p><b>Values</b> 1 to 4294967295   <i>te-link-name</i><br/> <i>te-link-name</i>: 32 character maximum name of the TE Link</p>                         |

## shutdown

|                    |                                                        |
|--------------------|--------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                   |
| <b>Context</b>     | config>router>gmpls>te-link                            |
| <b>Description</b> | This command enables or disables the TE Link in GMPLS. |
| <b>Default</b>     | no shutdown                                            |

---

### 3.19.2.2.1 GMPLS Tunnel Group Commands

#### gmpls-tun-grp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no]</b> <b>gmpls-tun-grp</b> <i>gmpls-tunnel-group-id</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command configures a GMPLS tunnel group. A GMPLS tunnel group is a bundle of GMPLS LSPs providing an abstraction of the data bearers that are intended to be associated to one IP interface. This object allows, for example, end-to-end load balancing across the set of data bearers corresponding to a set of gLSPs. A gLSP is bound to an overlay tunnel group by a gLSP tunnel name at both the head end and the tail end UNI-C nodes of a gLSP. A sender-address may be optionally configured for the tail end of a gLSP in case overlapping GMPLS LSP tunnel names are used by different head end nodes.</p> <p>The <b>no</b> form of this command removes the tunnel group. All members of a GMPLS tunnel group must be removed and the tunnel group shutdown before the tunnel group can be deleted.</p> |
| <b>Default</b>     | no gmpls-tun-grp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <i>gmpls-tunnel-group-id</i> — Specifies the identifier of the GMPLS tunnel group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Values</b>      | 1 to 1024                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

#### description

|                    |                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>description</b> <i>description-string</i><br><b>no description</b>                                                                              |
| <b>Context</b>     | config>gmpls-tun-grp                                                                                                                               |
| <b>Description</b> | <p>This command configures a description string for the GMPLS tunnel group.</p> <p>The <b>no</b> form of this command removes the description.</p> |
| <b>Default</b>     | no description                                                                                                                                     |
| <b>Parameters</b>  | <i>description-string</i> — Specifies a text string of up to 160 characters describing the GMPLS tunnel group.                                     |

#### far-end

|               |                                                       |
|---------------|-------------------------------------------------------|
| <b>Syntax</b> | <b>far-end</b> <i>ip-address</i><br><b>no far-end</b> |
|---------------|-------------------------------------------------------|

---

|                    |                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>gmpls-tun-grp                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command configures the IP address (GMPLS Loopback Address) of the far-end UNI-C router.</p> <p>The <b>no</b> form of this command removes the far-end address</p>                                                                                                      |
| <b>Default</b>     | no far-end                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <p><i>ip-address</i> — Specifies an IPv4 or IPv6 address of the far-end UNI-C router.</p> <p><b>Values</b></p> <p>ipv4-address — a.b.c.d</p> <p>ipv6-address — x:x:x:x:x:x:x (eight 16-bit pieces)</p> <p>x:x:x:x:x:d.d.d.d</p> <p>x — [0 to FFFF]H</p> <p>d — [0 to 255]D</p> |

## member

|                    |                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] member <i>member-id</i></b>                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>gmpls-tun-grp                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command configures a member of a GMPLS tunnel group. A member of a GMPLS tunnel group is a GMPLS LSP. All members of a tunnel group must have the same bandwidth. Up to 16 members may be configured for each GMPLS tunnel group.</p> <p>The <b>no</b> form of this command removes the member.</p> |
| <b>Default</b>     | no member                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <p><i>member-id</i> — Specifies the identifier of the GMPLS tunnel group member.</p> <p><b>Values</b> 1 to 16</p>                                                                                                                                                                                           |

## glsp

|                    |                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] glsp session-name <i>name</i></b>                                                                                                                                                                                             |
| <b>Context</b>     | config>gmpls-tun-grp>member                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command binds a GMPLS LSP as a member of the GMPLS tunnel group. The session name is used to identify the GMPLS LSP. This is the LSP name of the GMPLS LSP.</p> <p>The <b>no</b> form of this command removes the member.</p> |
| <b>Default</b>     | none                                                                                                                                                                                                                                  |

---

|                   |                                                                         |
|-------------------|-------------------------------------------------------------------------|
| <b>Parameters</b> | <b>session-name name</b> — Specifies the session name of the GMPLS LSP. |
| <b>Values</b>     | 80 characters maximum text string                                       |

shutdown

|                    |                                                                        |
|--------------------|------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                                   |
| <b>Context</b>     | config>gmpls-tun-grp>member                                            |
| <b>Description</b> | This command disables or enables the member of the GMPLS tunnel group. |
| <b>Default</b>     | shutdown                                                               |

member-threshold

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] member-threshold <i>threshold</i></b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>gmpls-tun-grp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>The <b>member-threshold</b> is the number of member GMPLS LSPs that must be operationally up before the GMPLS tunnel group is considered operationally up. If that number is not reached, then the GMPLS tunnel group is taken operationally down.</p> <p>A member of a GMPLS tunnel group may be treated as down for one of the following reasons. These reason codes are recorded in the tmnxGmplsTunGrpMemberTable in the MIB:</p> <ul style="list-style-type: none"><li>• adminDn — The member or the related tunnel group is admin down.</li><li>• wpLspDn — The associated GMPLS LSP working path is down.</li><li>• wpPortDn — The data-bearer port associated with the GMPLS LSP working path is down.</li><li>• wpPortNoRsrc — The data bearer port associated with the LSP working path has no resource to support the services over the GMPLS tunnel group logical port.</li><li>• ppLspDn — The associated GMPLS LSP protect path is down.</li><li>• ppPortDn — The data-bearer port associated with the GMPLS LSP protect path is down.</li><li>• ppPortNoRsrc — The data bearer port associated with the GMPLS LSP protect path has no resource to support the services over the GMPLS tunnel group logical port.</li></ul> <p>The <b>no</b> form of this command reverts the member threshold to 0.</p> |
| <b>Default</b>     | 0                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <b><i>threshold</i></b> — Specifies the minimum number of GMPLS LSPs that must be operationally up before the GMPLS tunnel group is considered operationally up.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Values</b>      | 0 to 15                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## mode

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |               |                                                                                                                                      |                |                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|--------------------------------------------------------------------------------------------------------------------------------------|----------------|---------------------|
| <b>Syntax</b>      | <b>mode</b> <i>mode</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |               |                                                                                                                                      |                |                     |
| <b>Context</b>     | config>gmpls-tun-grp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |               |                                                                                                                                      |                |                     |
| <b>Description</b> | <p>This command sets the operating mode of the GMPLS tunnel group.</p> <p>In <b>load-sharing</b> mode, traffic is load-shared across the member GMPLS LSPs of the tunnel group. The same hashing algorithm is used as for LAG (see the "LAG and ECMP hashing" chapter of the 7450 ESS, 7750 SR, and 7950 XRS Interface Configuration Guide). If load-sharing is configured, then all of the GMPLS LSPs must terminate on the same far-end node. All of the ports used by GMPLS LSPs must be equivalent in that they must have the same named QoS policy, bandwidth, etc. Once more than one gLSP is associated with a tunnel group, the QoS policy / scheduler policy cannot be changed for any of the ports. All GMPLS LSPs must be unprotected end-to-end. Segment protection is allowed for GMPLS LSPs associated in a load sharing mode tunnel group.</p> <p>In <b>active-standby</b> mode, only one member gLSP can be associated with the tunnel group.</p> <p>The <b>no</b> form of this command removes the member.</p> |               |                                                                                                                                      |                |                     |
| <b>Default</b>     | load-sharing                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |               |                                                                                                                                      |                |                     |
| <b>Parameters</b>  | <p><i>mode</i> — Specifies the operating mode of the GMPLS tunnel group.</p> <table> <tr> <td><b>Values</b></td><td> <b>active-standby</b> — Sets the operating mode to active-standby.<br/> <b>load-sharing</b> — Sets the operating mode to load-sharing. </td></tr> <tr> <td><b>Default</b></td><td><b>load-sharing</b></td></tr> </table>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>Values</b> | <b>active-standby</b> — Sets the operating mode to active-standby.<br><b>load-sharing</b> — Sets the operating mode to load-sharing. | <b>Default</b> | <b>load-sharing</b> |
| <b>Values</b>      | <b>active-standby</b> — Sets the operating mode to active-standby.<br><b>load-sharing</b> — Sets the operating mode to load-sharing.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |               |                                                                                                                                      |                |                     |
| <b>Default</b>     | <b>load-sharing</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |               |                                                                                                                                      |                |                     |

## shutdown

|                    |                                                                           |
|--------------------|---------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                                      |
| <b>Context</b>     | config>gmpls-tun-grp                                                      |
| <b>Description</b> | This command administratively disables or enables the GMPLS tunnel group. |
| <b>Default</b>     | shutdown                                                                  |

## type

|                |                                   |
|----------------|-----------------------------------|
| <b>Syntax</b>  | <b>type</b> [head-end   tail-end] |
| <b>Context</b> | config>gmpls-tun-grp              |

---

|                    |                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command configures whether a GMPLS tunnel group is at the head-end or tail-end of the set of member GMPLS LSPs from the perspective of GMPLS LSP setup. It can only be configured if the GMPLS tunnel group has no members; for example, if none have yet been configured. |
| <b>Default</b>     | head-end                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <b>head-end</b> — Sets the GMPLS tunnel group to operate as a head-end.<br><b>tail-end</b> — Sets the GMPLS tunnel group to operate as a tail-end.                                                                                                                              |

## 3.20 Show, Clear, Debug, and Tools Command Reference

### 3.20.1 Command Hierarchies

- [Show Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)
- [Tools Commands](#)

#### 3.20.1.1 Show Commands

```
show
 — router
 — gmpls
 — lsp [lsp-name] [status {up | down}] [detail]
 — lsp [lsp-name] path [path-name] [type {working | protect}] [detail]
 — lsp [lsp-name] [status {up | down}] to ip-address [detail]
 — peer [peer-node-id] [detail]
 — peer [peer-node-id] {detail}
 — path [path-name]
 — session [session-type] [from ip-address] [to ip-address] [session-name
 session-name] [status {up | down}] [detail]
 — te-link [te-link-id]
 — Imp
 — peer [peer-node-id] [detail]
 — control-channel Imp-cc-id [detail]
 — te-link [te-link-id] [detail]
 — te-link te-link-id statistics
 — data-bearer [data-bearer-id] [detail]
 — gmpls-tun-grp [gmpls-tunnel-group-id] [using]
 — member [member-id]
```

#### 3.20.1.2 Clear Commands

```
clear
 — router
 — statistics
```

### 3.20.1.3 Debug Commands

```

debug
 — router
 — gmpls [lsp lsp-name] [sender sender-address] [endpoint endpoint-address]
 [tunnel-id tunnel-id] [lsp-id lsp-id] [peer peer-node-id]
 — no gmpls
 — [no] event
 — [no] lsp-setup
 — [no] path
 — [no] resv
 — [no] peer
 — [no] rr
 — [no] misc
 — [no] all
 — [no] packet
 — hello [detail]
 — no hello
 — path [detail]
 — no path
 — patherr [detail]
 — no patherr
 — pathtear [detail]
 — no pathtear
 — resv [detail]
 — no resv
 — resverr [detail]
 — no resverr
 — resvtear [detail]
 — no resvtear
 — notify [detail]
 — no notify
 — ack [detail]
 — no ack
 — srefresh [detail]
 — no srefresh
 — bundle [detail]
 — no bundle
 — all [detail]
 — no all

```

### 3.20.1.4 Tools Commands

```

tools
 — dump
 — gmpls-tun-grp id
 — router
 — lmp
 — peer peer-node-id
 — control-channel lmp-cc-id

```



- **te-link** *te-link-id*
- **gmpls**
  - **lsp** [*lsp-name*] [detail]
  - **path** [*path-name*] [detail]
  - **peer** [*peer-node-id*] [detail]
  - **port** [*session-name*] [*sender ip-address*] [detail]
  - **psb** [*tunnelid tunnel-id*] [*lspid lsp-id*] [detail]
  - **status**
  - **te-link** [*te-link-id*] [detail]

## 3.20.2 Command Descriptions

### 3.20.2.1 Show Commands



**Note:** The command outputs in this chapter are examples only; actual displays may differ depending on supported functionality and user configuration.

#### gmpls

|                    |                                                                               |
|--------------------|-------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>gmpls</b>                                                                  |
| <b>Context</b>     | show>router>gmpls                                                             |
| <b>Description</b> | This command displays RSVP status information for the GMPLS instance of RSVP. |

#### lsp

|                    |                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lsp</b> [ <i>lsp-name</i> ] [status {up   down}] [detail]<br><b>lsp</b> [ <i>lsp-name</i> ] path [ <i>path-name</i> ] [type {working   protect}] [detail]<br><b>lsp</b> [ <i>lsp-name</i> ] [status {up   down}] to <i>ip-address</i> [detail]                                                                                                                                     |
| <b>Context</b>     | show>router>gmpls                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | This command displays gLSP information.                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <i>lsp-name</i> — Specifies the identifier for the GMPLS LSP. The LSP name can be up to 32 characters long and must be unique.<br><b>status</b> — Specifies that the command display only LSPs that are either operationally up or down.<br><b>Values</b><br><b>up</b> — Display only LSPs that are operationally up.<br><b>down</b> — Display only LSPs that are operationally down. |

*path-name* — Specifies a unique case-sensitive name label for the LSP path.

**Values** 32 characters maximum alphanumeric string

**type** — Specifies that the command display only paths that are either working or protect.

**Values** **working** — Display only working paths.

**protect** — Display only protect paths.

**to ip-address** — Specifies the system IP address of the far-end UNI-C router for the GMPLS LSP.

**detail** — Keyword to request more detailed output.

## peer

**Syntax** **peer** [*peer-node-id*] [**detail**]  
**peer** [*peer-node-id*] {**statistics**}

**Context** show>router>gmpls

**Description** This command displays GMPLS peer information.

**Parameters** *peer-node-id* — Specifies the control plane node ID of the neighboring GMPLS UNI-N node. This can be an IP address or a 32-bit unsigned integer.

**Values** {a.b.c.d | 1 — 4294967295}

**detail** — Keyword to request more detailed output.

**statistics** — Keyword to request peer statistics.

## path

**Syntax** **path** [*path-name*]

**Context** show>router>gmpls

**Description** This command displays GMPLS path information.

**Parameters** *path-name* — Specifies a unique case-sensitive name label for the LSP path.

**Values** 32 characters maximum alphanumeric string

## session

**Syntax** **session** [**session-type**] [**from ip-address**] [**to ip-address**] [**session-name session-name**]  
[**status {up | down}**] [**detail**]

**Context** show>router>gmpls

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command displays GMPLS session information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><b>session-type</b> — Keyword to display information about the session type.</p> <p><b>from</b> <i>ip-address</i> — Specifies the system IP address of the near-end UNI-C router.</p> <p><b>to</b> <i>ip-address</i> — Specifies the system IP address of the far-end UNI-C router.</p> <p><i>session-name</i> — Specifies the name of the GMPLS session.</p> <p><b>status</b> — Specifies that the command display only GMPLS sessions that are either operationally up or down.</p> <p><b>Values</b></p> <p><b>up</b> — Display only sessions that are operationally up.</p> <p><b>down</b> — Display only sessions that are operationally down.</p> <p><b>detail</b> — Keyword to request more detailed output.</p> |

## te-link

|                    |                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>te-link</b> [ <i>te-link-id</i> ]                                                                                                                                                                                 |
| <b>Context</b>     | show>router>gmpls                                                                                                                                                                                                    |
| <b>Description</b> | This command displays Traffic Engineering (TE) link information.                                                                                                                                                     |
| <b>Parameters</b>  | <p><i>te-link-id</i> — Specifies the ID or name of the configured TE Link.</p> <p><b>Values</b></p> <p>1 — 4294967295   <i>te-link-name</i></p> <p><i>te-link-name</i>: 32 character maximum name of the TE Link</p> |

## gmpls-tun-grp

|                    |                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>gmpls-tun-grp</b> [ <i>gmpls-tunnel-group-id</i> ] [ <b>using</b> ]                                                                                                                                                     |
| <b>Context</b>     | show                                                                                                                                                                                                                       |
| <b>Description</b> | This command displays GMPLS tunnel group status and lists the gLSPs bound to each tunnel group.                                                                                                                            |
| <b>Parameters</b>  | <p><b>id</b> <i>group-id</i> — Specifies that only GMPLS tunnel groups of the configured ID are displayed.</p> <p><b>using</b> — Keyword to display information about which IP interfaces are using the tunnel groups.</p> |

## member

|               |                                    |
|---------------|------------------------------------|
| <b>Syntax</b> | <b>member</b> [ <i>member-id</i> ] |
|---------------|------------------------------------|

|                    |                                                                                |
|--------------------|--------------------------------------------------------------------------------|
| <b>Context</b>     | show>gmpls-tun-grp gmpls-tunnel-group-id                                       |
| <b>Description</b> | This command displays member information for the specified GMPLS tunnel group. |
| <b>Parameters</b>  | <i>member-id</i> — Specifies the ID of the GMPLS tunnel group member.          |
| <b>Values</b>      | 1 — 16                                                                         |

## peer

|                    |                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>peer</b> [ <i>peer-node-id</i> ] [ <b>detail</b> ]                                                                                    |
| <b>Context</b>     | show>router>Imp                                                                                                                          |
| <b>Description</b> | This command displays LMP peer information.                                                                                              |
| <b>Parameters</b>  | <i>peer-node-id</i> — Specifies the unique identifier for the LMP peer node.<br><b>detail</b> — Keyword to display more detailed output. |
| <b>Output</b>      | Peer Output Fields<br><br><a href="#">Table 30</a> describes the peer output fields.                                                     |

**Table 30** Peer Output Fields

| Label         | Description                                                                         |
|---------------|-------------------------------------------------------------------------------------|
| Admin State   | The administrative state of the LMP peer node.                                      |
| Oper State    | The operational state of the LMP peer node.                                         |
| Retrans Intvl | The configured interval between resubmitted LMP messages.                           |
| Retry Limit   | The configured number of times LMP resends a message before restarting the process. |

### Sample Output

```
A:ALA-1# show router lmp peer 16843009
=====
LMP Peer 1.1.1.1 (16843009)
=====
Admin State: : Out of Service Oper State : Down
Retrans Intvl : 5000 ms Retry Limit : 6
=====
A:ALA-1#
```

## control-channel

- Syntax** `control-channel Imp-cc-id [detail]`
- Context** `show>router>Imp>peer`
- Description** This command displays LMP control channel information for a specific peer.
- Parameters** *Imp-cc-id* — Specifies the unsigned integer identifier for the control channel
- Values** 1 — 42949672
- detail** — Keyword to display more detailed output.
- Output** Control Channel Output Fields

[Table 31](#) describes control channel output fields.

**Table 31 Control Channel Output Fields**

| Label           | Description                                                                                            |
|-----------------|--------------------------------------------------------------------------------------------------------|
| Admin State     | The administrative state of the control channel.                                                       |
| Oper State      | The operational state of the control channel.                                                          |
| Remote Id       | The remote ID of the control channel.                                                                  |
| Setup Role      | The setup role of the control channel.                                                                 |
| Hello Interval  | The interval at which LMP hello packets are sent on the control channel.                               |
| Hello Dead Intv | The interval after which the IPCC is declared down if no hello packets are received from the LMP peer. |
| Peer If Addr    | The LMP peer interface address.                                                                        |
| Inbound Errors  | The total number of inbound packet errors.                                                             |
| Outbound Errors | The total number of outbound packet errors.                                                            |
| Message Type    | The types of message packets sent and received on the control channel.                                 |
| Received        | The total number of received packets of a specific type.                                               |
| Sent            | The total number of sent packets of a specific type.                                                   |
| Retransmitted   | The total number of retransmitted packets of a specific type.                                          |

### Sample Output

```
A:ALA-1# show router lmp peer 16843009 control-channel 1
```

```
=====
```

```

LMP Control Channel 1 Statistics
=====
Admin State: : Out of Service Oper State : Going Down
Remote Id : 10 Setup Role : Active
Hello Interval : 3000 Hello Dead Intv : 10000
Peer If Addr : 100.100.100.100
Inbound Errors : 100 Outbound Errors : 50

Message Type Received Sent Retransmitted

Config 1111 11 111
ConfigAck 15 6 N/A
ConfigNack 1 2 N/A
Hello 40 20 N/A
LinkSummary 6000 500 10
LinkSummaryAck 35 30 N/A
LinkSummaryNack 400 3000 N/A
=====
A:ALA-1#

```

## te-link

- Syntax** **te-link** [*te-link-id*] [**detail**]  
**te-link** *te-link-id* **statistics**
- Context** show>router>Imp
- Description** This command displays Traffic Engineering (TE) link information.
- Parameters** *te-link-id* — Specifies the ID of a TE Link.
- Values** 1 — 4294967295 | *te-link-name*  
*te-link-name*: 32 character (max) name of the TE Link
- detail** — Keyword to display more detailed output.
- statistics** — Keyword to display TE Link statistics.
- Output** TE Link Output Fields
- [Table 32](#) describes TE link output fields.

**Table 32** TE Link Output Fields

| Label        | Description                              |
|--------------|------------------------------------------|
| Admin State  | The administrative state of the TE link. |
| Oper State   | The operational state of the TE link.    |
| Remote Id    | The remote ID of the TE link.            |
| Peer Node Id | The LMP peer node ID.                    |

**Table 32** TE Link Output Fields (Continued)

| Label         | Description (Continued)                                       |
|---------------|---------------------------------------------------------------|
| Received      | The total number of received packets of a specific type.      |
| Sent          | The total number of sent packets of a specific type.          |
| Retransmitted | The total number of retransmitted packets of a specific type. |

**Sample Output**

```
*A:SRU4>show>router>lmp# te-link 254
=====
LMP TE Link 254 (Name) Statistics
=====
Admin State : In Service Oper State : Degraded
Remote Id : 10 Peer Node Id : 15

Message Type Received Sent Retransmitted

LinkSummary 100 100 600
LinkSummaryAck 30 25 N/A
LinkSummaryNack 70 75 N/A
=====
```

**data-bearer**

- Syntax** **data-bearer** [*data-bearer-id*] [**detail**]
- Context** show>router>lmp>te-link
- Description** This command displays LMP data bearer information.
- Parameters** *data-bearer-id* — Specifies the ID of a data bearer link.  
**detail** — Keyword to display more detailed output.
- Output** Data Bearer Output Fields

[Table 33](#) describes the data bearer output fields.

**Table 33** Data Bearer Output Fields

| Label       | Description                                       |
|-------------|---------------------------------------------------|
| Admin State | The administrative state of the data bearer link. |
| Oper State  | The operational state of the data bearer link.    |
| Remote Id   | The remote ID of the data bearer link.            |
| Port        | The port used by the data bearer link.            |

**Sample Output**

```
*A:SRU4>show>router>lmp>te-link# data-bearer 254
=====
LMP Data Bearer Link 254 Statistics
=====
Admin State : In Service Oper State : Up Free
Remote Id : 10 Port : 1/1/1
=====
```

### 3.20.2.2 Clear Commands

statistics

|                    |                                                |
|--------------------|------------------------------------------------|
| <b>Syntax</b>      | <b>statistics</b>                              |
| <b>Context</b>     | clear>router>gmpls>peer                        |
| <b>Description</b> | This command clears control packet statistics. |

### 3.20.2.3 Debug Commands

gmpls

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>gmpls [lsp <i>lsp-name</i>] [sender <i>sender-address</i>] [endpoint <i>endpoint-address</i>] [tunnel-id <i>tunnel-id</i>] [lsp-id <i>lsp-id</i>] [peer <i>peer-node-id</i>]<br/>no gmpls</b>                                                                                                                                                                                                                                   |
| <b>Context</b>     | debug>router                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | This command enables and configures debugging for GMPLS.<br><br>The <b>no</b> form of the command disables debugging for GMPLS.                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b>  | <i>lsp-name</i> — Debug all GMPLS instances that contain the specified LSP. 80 characters max.<br><br><i>sender-address</i> — Debug all GMPLS instances that contain the specified sender IP address, in dotted decimal notation.<br><br><b>Values</b> a.b.c.d<br><br><i>endpoint-address</i> — Debug all GMPLS instances that contain the specified endpoint IP address, in dotted decimal notation.<br><br><b>Values</b> a.b.c.d |



*tunnel-id* — Debug all GMPLS instances that contain the specified identifier of a tunnel.

**Values** 0 — 4294967295

*lsp-id* — Debug all GMPLS instances that contain the specified identifier of an LSP.

**Values** 1 — 65535

*peer-node-id* — Debug all GMPLS instances that contain the specified identifier or IP address of a peer node.

**Values** 1 — 4294967295 | a.b.c.d

## event

**Syntax** [no] event

**Context** debug>router>gmpls

**Description** This command enables and disables debugging for specific GMPLS events.

## all

**Syntax** [no] all

**Context** debug>router>gmpls

**Description** This command enables and disables debugging for GMPLS All events.

## lsp-setup

**Syntax** [no] lsp-setup

**Context** debug>router>gmpls

**Description** This command enables and disables debugging for GMPLS LSP Setup events.

## misc

**Syntax** [no] misc

**Context** debug>router>gmpls

**Description** This command enables and disables debugging for GMPLS Misc events.

## path

|                    |                                                                    |
|--------------------|--------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] path</b>                                                   |
| <b>Context</b>     | debug>router>gmpls                                                 |
| <b>Description</b> | This command enables and disables debugging for GMPLS Path events. |

## peer

|                    |                                                                   |
|--------------------|-------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] peer</b>                                                  |
| <b>Context</b>     | debug>router>gmpls                                                |
| <b>Description</b> | This command enables and disables debugging for GMPLS NBR events. |

## resv

|                    |                                                                    |
|--------------------|--------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] resv</b>                                                   |
| <b>Context</b>     | debug>router>gmpls                                                 |
| <b>Description</b> | This command enables and disables debugging for GMPLS Resv events. |

## rr

|                    |                                                                                 |
|--------------------|---------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] rr</b>                                                                  |
| <b>Context</b>     | debug>router>gmpls                                                              |
| <b>Description</b> | This command enables and disables debugging for GMPLS Refresh Reduction events. |

## packet

|                    |                                                                         |
|--------------------|-------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] packet</b>                                                      |
| <b>Context</b>     | debug>router>gmpls                                                      |
| <b>Description</b> | This command enables and disables debugging for specific GMPLS packets. |

## hello

|               |                                    |
|---------------|------------------------------------|
| <b>Syntax</b> | <b>hello [detail]<br/>no hello</b> |
|---------------|------------------------------------|

---

|                    |                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | debug>router>gmpls>packet                                                                                                                    |
| <b>Description</b> | This command enables debugging for GMPLS Hello packets.<br><br>The <b>no</b> form of the command disables debugging for GMPLS Hello packets. |
| <b>Parameters</b>  | <b>detail</b> — Keyword to produce debug output in greater detail.                                                                           |

## path

|                    |                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>path [detail]</b><br><b>no path</b>                                                                                                     |
| <b>Context</b>     | debug>router>gmpls>packet                                                                                                                  |
| <b>Description</b> | This command enables debugging for GMPLS Path packets.<br><br>The <b>no</b> form of the command disables debugging for GMPLS Path packets. |
| <b>Parameters</b>  | <b>detail</b> — Keyword to produce debug output in greater detail.                                                                         |

## patherr

|                    |                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>patherr [detail]</b><br><b>no patherr</b>                                                                                                     |
| <b>Context</b>     | debug>router>gmpls>packet                                                                                                                        |
| <b>Description</b> | This command enables debugging for GMPLS PathErr packets.<br><br>The <b>no</b> form of the command disables debugging for GMPLS PathErr packets. |
| <b>Parameters</b>  | <b>detail</b> — Keyword to produce debug output in greater detail.                                                                               |

## pathtear

|                    |                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>pathtear [detail]</b><br><b>no pathtear</b>                                                                                                     |
| <b>Context</b>     | debug>router>gmpls>packet                                                                                                                          |
| <b>Description</b> | This command enables debugging for GMPLS PathTear packets.<br><br>The <b>no</b> form of the command disables debugging for GMPLS PathTear packets. |
| <b>Parameters</b>  | <b>detail</b> — Keyword to produce debug output in greater detail.                                                                                 |

## resv

|                    |                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>resv [detail]</b><br><b>no resv</b>                                                                                                     |
| <b>Context</b>     | debug>router>gmpls>packet                                                                                                                  |
| <b>Description</b> | This command enables debugging for GMPLS Resv packets.<br><br>The <b>no</b> form of the command disables debugging for GMPLS Resv packets. |
| <b>Parameters</b>  | <b>detail</b> — Keyword to produce debug output in greater detail.                                                                         |

## resvrr

|                    |                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>resvrr [detail]</b><br><b>no resvrr</b>                                                                                                       |
| <b>Context</b>     | debug>router>gmpls>packet                                                                                                                        |
| <b>Description</b> | This command enables debugging for GMPLS ResvErr packets.<br><br>The <b>no</b> form of the command disables debugging for GMPLS ResvErr packets. |
| <b>Parameters</b>  | <b>detail</b> — Keyword to produce debug output in greater detail.                                                                               |

## resvtear

|                    |                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>resvtear [detail]</b><br><b>no resvtear</b>                                                                                                     |
| <b>Context</b>     | debug>router>gmpls>packet                                                                                                                          |
| <b>Description</b> | This command enables debugging for GMPLS ResvTear packets.<br><br>The <b>no</b> form of the command disables debugging for GMPLS ResvTear packets. |
| <b>Parameters</b>  | <b>detail</b> — Keyword to produce debug output in greater detail.                                                                                 |

## notify

|                    |                                                          |
|--------------------|----------------------------------------------------------|
| <b>Syntax</b>      | <b>notify [detail]</b><br><b>no notify</b>               |
| <b>Context</b>     | debug>router>gmpls>packet                                |
| <b>Description</b> | This command enables debugging for GMPLS Notify packets. |

The **no** form of the command disables debugging for GMPLS Notify packets.

**Parameters**     **detail** — Keyword to produce debug output in greater detail.

## ack

**Syntax**     **ack [detail]**  
**no ack**

**Context**     debug>router>gmpls>packet

**Description**     This command enables debugging for GMPLS Ack packets.

The **no** form of the command disables debugging for GMPLS Ack packets.

**Parameters**     **detail** — Keyword to produce debug output in greater detail.

## srefresh

**Syntax**     **srefresh [detail]**  
**no srefresh**

**Context**     debug>router>gmpls>packet

**Description**     This command enables debugging for GMPLS Srefresh packets.

The **no** form of the command disables debugging for GMPLS Srefresh packets.

**Parameters**     **detail** — Keyword to produce debug output in greater detail.

## bundle

**Syntax**     **bundle [detail]**  
**no bundle**

**Context**     debug>router>gmpls>packet

**Description**     This command enables debugging for GMPLS Bundle packets.

The **no** form of the command disables debugging for GMPLS Bundle packets.

**Parameters**     **detail** — Keyword to produce debug output in greater detail.

## all

**Syntax**     **all [detail]**

**no all**

|                    |                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | debug>router>gmpls>packet                                                                                                                |
| <b>Description</b> | This command enables debugging for GMPLS All packets.<br><br>The <b>no</b> form of the command disables debugging for GMPLS All packets. |
| <b>Parameters</b>  | <b>detail</b> — Keyword to produce debug output in greater detail.                                                                       |

### 3.20.2.4 Tools Commands

#### gmpls-tun-grp

|                    |                                                                          |
|--------------------|--------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>gmpls-tun-grp <i>gmpls-tunnel-group-id</i></b>                        |
| <b>Context</b>     | tools>dump                                                               |
| <b>Description</b> | This command dumps information about a GMPLS tunnel group.               |
| <b>Parameters</b>  | <i>gmpls-tunnel-group-id</i> — The identifier of the GMPLS tunnel group. |
| <b>Values</b>      | 1 — 1024                                                                 |

#### Output

#### Sample Output

```
*A:Dut-A-UNI-C# tools dump gmpls-tun-grp 1

TunnelGrpId 1: (Up/Up), Port gmpls-tun-grp-1 (Up/Up), mode load-sharing
Type: Head, far-end: 52.52.52.52, bw 40000000, memThreshold 0, portId 0x5c000001
NumMems: 5/4/0(0), Up/Dn 1/0, active 0x1e, present 0x1e iom 0x2 Mtu 9212
gmplsUpd: 8 (000 01:11:23.350), Now: 000 08:34:12.970
 memId 1: (Up/Dn), session: 1::1, reg: Y
 DnReasons: wpLspDn
 Work: N/A (Ghost), status: Dn, Cnt(Up/Dn/PortChg): 0/0/0
 lspState Dn, lspUpd: 0 (000 00:00:00.000), rsrcAdded N
 memId 2: (Up/Up), session: 1::empty, reg: Y
 Work: 1/2/1 (Up), status: Up, Cnt(Up/Dn/PortChg): 1/0/1
 lspState Up, lspUpd: 1 (000 01:10:14.720), rsrcAdded Y
 memId 3: (Up/Up), session: 2::empty, reg: Y
 Work: 1/2/2 (Up), status: Up, Cnt(Up/Dn/PortChg): 1/0/1
 lspState Up, lspUpd: 1 (000 01:10:23.650), rsrcAdded Y
 memId 4: (Up/Up), session: 3::empty, reg: Y
 Work: 1/1/14 (Up), status: Up, Cnt(Up/Dn/PortChg): 1/0/1
 lspState Up, lspUpd: 1 (000 01:10:05.880), rsrcAdded Y
 memId 5: (Up/Up), session: 4::empty, reg: Y
 Work: 1/1/13 (Up), status: Up, Cnt(Up/Dn/PortChg): 1/0/1
 lspState Up, lspUpd: 1 (000 01:09:50.710), rsrcAdded Y

Sorted list of 4 member port(s):
```

```
1/1/13 1/1/14 1/2/1 1/2/2

*A:Dut-A-UNI-C# tools dump gmpls-tun-grp 1
clear

*A:Dut-A-UNI-C# tools dump gmpls-tun-grp 1 clear

TunnelGrpId 1: (Up/Up), Port gmpls-tun-grp-1 (Up/Up), mode load-sharing
Type: Head, far-end: 52.52.52.52, bw 40000000, memThreshold 0, portId 0x5c000001
NumMems: 5/4/0(0), Up/Dn 1/0, active 0x1e, present 0x1e iom 0x2 Mtu 9212
gmplsUpd: 8 (000 01:11:23.350), Now: 000 08:34:20.390
 memId 1: (Up/Dn), session: 1::1, reg: Y
 DnReasons: wpLspDn
 Work: N/A (Ghost), status: Dn, Cnt(Up/Dn/PortChg): 0/0/0
 lspState Dn, lspUpd: 0 (000 00:00:00.000), rsrcAdded N
 memId 2: (Up/Up), session: 1::empty, reg: Y
 Work: 1/2/1 (Up), status: Up, Cnt(Up/Dn/PortChg): 1/0/1
 lspState Up, lspUpd: 1 (000 01:10:14.720), rsrcAdded Y
 memId 3: (Up/Up), session: 2::empty, reg: Y
 Work: 1/2/2 (Up), status: Up, Cnt(Up/Dn/PortChg): 1/0/1
 lspState Up, lspUpd: 1 (000 01:10:23.650), rsrcAdded Y
 memId 4: (Up/Up), session: 3::empty, reg: Y
 Work: 1/1/14 (Up), status: Up, Cnt(Up/Dn/PortChg): 1/0/1
 lspState Up, lspUpd: 1 (000 01:10:05.880), rsrcAdded Y
 memId 5: (Up/Up), session: 4::empty, reg: Y
 Work: 1/1/13 (Up), status: Up, Cnt(Up/Dn/PortChg): 1/0/1
 lspState Up, lspUpd: 1 (000 01:09:50.710), rsrcAdded Y

Sorted list of 4 member port(s):
1/1/13 1/1/14 1/2/1 1/2/2
```

## peer

|                    |                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>peer</b> <i>peer-node-id</i>                                                                                                |
| <b>Context</b>     | tools>dump>router>Imp                                                                                                          |
| <b>Description</b> | This command dumps information about a specific LMP peer.                                                                      |
| <b>Parameters</b>  | <i>peer-node-id</i> — An identifier for the LMP peer node. This may be an IPv4-formatted address or a 32-bit unsigned integer. |
| <b>Values</b>      | a.b.c.d   1 — 4294967295                                                                                                       |

## control-channel

|                    |                                                                        |
|--------------------|------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>control-channel</b> <i>Imp-cc-id</i>                                |
| <b>Context</b>     | tools>dump>router>Imp>peer                                             |
| <b>Description</b> | This command dumps information about a control channel of an LMP peer. |

**Parameters**    *lmp-cc-id* — An unsigned integer identifier for the control channel.

**Values**        1 — 42949672

## Output

### Sample Output

```
*A:Dut-B-UNI-C>tools>dump>router>lmp# peer 10.27.255.21
10.27.255.213 10.27.255.219
*A:Dut-B-UNI-C>tools>dump>router>lmp# peer 10.27.255.21 control-channel 7001
MINOR: CLI peer 10.27.255.21 is not configured.
*A:Dut-B-UNI-C>tools>dump>router>lmp# peer 10.27.255.21
10.27.255.213 10.27.255.219
*A:Dut-B-UNI-C>tools>dump>router>lmp# peer 10.27.255.213 control-channel
<lmp-cc-id>
7001
```

```
*A:Dut-B-UNI-C>tools>dump>router>lmp# peer 10.27.255.213 control-channel 7001
```

```
Local: 7001, Remote: 7004, State Up/Up, Nbr 10.27.255.213
Dst: 10.27.255.213, Nh: 172.21.36.173, EgrIfIndex: 3, Src: 52.52.52.52
Hello 1000/1000, Dead 3000/3000, Seq (Tx/Rx/Ack) 57995/16109/57995
LastEvt: rxHello, Tmrs: Hello Hold
Cnttrs: Errs(in/out) 0/0, Rx: cfg/Ack/Nack/Hello: 6/14/0/57887
Cnttrs: Tx: cfg/Ack/Nack/Hello: 14(30)/0/0/57995, reTxCnt: 0
RxCfg: 7004: 1000/3000, RxNack: 0: 0/0
TxMsgId: 44, RxMsgId(cfg/ack/nack): 1/44/0
LkSum Cnttrs (sum/ack/nack) Tx: 11654(216)/11620/0, Rx: 11724/11858/0
Now: 000 16:19:13.900
```

| Seq | Event   | prevState | state  | TxPkt | Time             |
|-----|---------|-----------|--------|-------|------------------|
| 009 | rxHello | Active    | Up     | ---   | 000 09:40:02.410 |
| 010 | holdTmr | Up        | CfgSnd | Cfg   | 000 10:15:14.410 |
| 011 | reTxTmr | CfgSnd    | CfgSnd | Cfg   | 000 10:15:14.900 |
| 012 | rxAck   | CfgSnd    | Active | Hello | 000 10:15:23.040 |
| 013 | rxHello | Active    | Up     | ---   | 000 10:15:24.020 |
| 014 | holdTmr | Up        | CfgSnd | Cfg   | 000 10:35:15.070 |
| 015 | rxAck   | CfgSnd    | Active | Hello | 000 10:35:15.080 |
| 016 | rxHello | Active    | Up     | ---   | 000 10:35:16.060 |
| 017 | holdTmr | Up        | CfgSnd | Cfg   | 000 10:45:20.090 |
| 018 | rxAck   | CfgSnd    | Active | Hello | 000 10:45:20.100 |
| 019 | rxHello | Active    | Up     | ---   | 000 10:45:21.080 |
| 000 | holdTmr | Up        | CfgSnd | Cfg   | 000 10:50:21.480 |
| 001 | rxAck   | CfgSnd    | Active | Hello | 000 10:50:21.490 |
| 002 | rxHello | Active    | Up     | ---   | 000 10:50:22.480 |
| 003 | holdTmr | Up        | CfgSnd | Cfg   | 000 11:05:23.550 |
| 004 | rxAck   | CfgSnd    | Active | Hello | 000 11:05:23.560 |
| 005 | rxHello | Active    | Up     | ---   | 000 11:05:24.540 |
| 006 | holdTmr | Up        | CfgSnd | Cfg   | 000 11:40:33.600 |
| 007 | rxAck   | CfgSnd    | Active | Hello | 000 11:40:33.610 |
| 008 | rxHello | Active    | Up     | ---   | 000 11:40:34.600 |



# te-link

|                    |                                                    |
|--------------------|----------------------------------------------------|
| <b>Syntax</b>      | <b>te-link</b> <i>te-link-id</i>                   |
| <b>Context</b>     | tools>dump>router>lmp>peer                         |
| <b>Description</b> | This command dumps information about a TE link.    |
| <b>Parameters</b>  | <i>te-link-id</i> — The identifier of the TE link. |

**Values** 1 — 4294967295 | *te-link-name*

## Output

### Sample Output

```
*A:Dut-B-UNI-C>tools>dump>router>lmp# te-link 1010701

Local 1010701 (), State Up/Up, Remote: 1010701, db 1(1), resv 0, msgId 11871
 Cntrs (sum/ack/nack) Tx: 2912(8)/2906/0, Rx: 2906/2914/0, Up/Dn: 14/13
 LastEvt: ackRxSm (000 16:19:18.070), Now: 000 16:19:33.270
 LastRxLkSum: Id 1010701, dbCnt 1, ErrCode 0x0, Flags: 0x0
 RxDb: Id(L/R) 101070117/101070117, Flags: 0x3
 SubObj: type/len/sw/enc: 1/12/125/2, Kbps(Min/Max): 10000000/10000000
 LastRxLkSumNack: msgId 0, dbCnt 0, ErrCode 0x0, Flags: 0x0
 Db Links:
 1010701-101070117 (1/1/13): State Up/UpAlloc, port Up/Up RemoteId: 101070117,
 Alloc Y, bw 10000000, portBw 10000000/10000000, tunGrp 1, cfg: Y, Lpc: Y
 Alloc (Cnt/Tm): 3/000 10:16:21.770, Free 2/000 10:15:32.050
 Cnt(Up/Dn): 1/0, UpDnTm: 000 00:10:55.310
```

# lsp

|                    |                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>lsp</b> [ <i>lsp-name</i> ] [ <b>detail</b> ]                                                                                         |
| <b>Context</b>     | tools>dump>router>gmpls                                                                                                                  |
| <b>Description</b> | This command dumps information about LSPs. Configuring the <i>lsp-name</i> parameter will only dump information about the specified LSP. |
| <b>Parameters</b>  | <i>lsp-name</i> — The name of the LSP. 32 characters max.<br><b>detail</b> — Keyword to output information in greater detail.            |

## Output

### Sample Output

```
*A:Dut-B-UNI-C>tools>dump>router>gmpls# lsp "D1"
LSP "D1", LspIndex: 1, Type: GMPLS-
UNI, AdminState: UP, OperState: UP, Flags: 0x0
Total LSP Count: 1
```

```
*A:Dut-B-UNI-C>tools>dump>router>gmpls# lsp "D1" detail
=====
LSP "D1", LspIndex: 1
 Type: GMPLS-UNI, AdminState: UP, OperState: UP, Flags: 0x0
 LastChange: 02/04/2015 23:30:34
 TimeSinceLastOperChange: 0d 15:07:19
 NumWorkingPaths: 1, NumProtectPaths: 0
 NumWorkingPathsUp: 1, NumProtectPathsUp: 0

LspPath PathIndex: 2, Type: Working
 AdminState: UP, OperState:Up, Flags: 0x0
 LastChange: 02/04/2015 23:30:34
 TimeSinceLastOperChange: 0d 15:07:19
 LspId: 296, SessionName: D1:2
 CurrState: LSP_PATH_UP, PrevState: LSP_PATH_INIT
 RetryCount: 107, FailNode: 0.0.0.0, FailCode: none
 Peer: 0.0.0.0
 PSB: vrId: 1 Session (To: 72.72.72.72 - 1 - 52.52.52.52), Sender (52.52.52.52 -
296)
 Actual Hops:
 Hop 1 : Unnumbered RouterId 52.52.52.52, InterfaceId 1010702
 UpStreamLabel 0, DnStreamLabel 101070217, Flags 0x0
 Hop 2 : Unnumbered RouterId 10.27.255.213, InterfaceId 1010702
 UpStreamLabel 101070217, DnStreamLabel 0, Flags 0x0
 SRLG List 27
 Hop 3 : Unnumbered RouterId 10.27.255.213, InterfaceId 13000
 UpStreamLabel 0, DnStreamLabel 101070230, Flags 0x0
 SRLG List 4
 Hop 4 : Unnumbered RouterId 10.27.255.219, InterfaceId 13001
 UpStreamLabel 101070230, DnStreamLabel 0, Flags 0x0
 Hop 5 : Unnumbered RouterId 10.27.255.219, InterfaceId 1010702
 UpStreamLabel 0, DnStreamLabel 101070217, Flags 0x0
 SRLG List 27
 Hop 6 : Unnumbered RouterId 72.72.72.72, InterfaceId 1010702
 UpStreamLabel 101070217, DnStreamLabel 0, Flags 0x0
=====
Total LSP Count: 1
```

path

|                    |                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>path</b> [ <i>path-name</i> ] [ <i>detail</i> ]                                                                                          |
| <b>Context</b>     | tools>dump>router>gmpls                                                                                                                     |
| <b>Description</b> | This command dumps information about paths. Configuring the <i>path-name</i> parameter will only dump information about the specified path. |
| <b>Parameters</b>  | <i>path-name</i> — The name of the path. 32 characters max.<br><b>detail</b> — Keyword to output information in greater detail.             |

Output

Sample Output

```
*A:Dut-B-UNI-C>tools>dump>router>gmpls# path
<path-name>
"empty" "2" "3" "5"
detail

*A:Dut-B-UNI-C>tools>dump>router>gmpls# path "2"
Path "2", PathIndex: 2, NumHops: 2, LspPathCount: 1, EmptyPathLspBindCount: 0,
Flags: 0x0
Total Path Count: 1

*A:Dut-B-UNI-C>tools>dump>router>gmpls# path "2" detail

Path "2", PathIndex: 2
AdminState: UP, OperState:UP
LastChange: 02/04/2015 23:30:34
NumHops: 2, LspPathCount: 1, EmptyPathLspBindCount: 0, Flags: 0x0
Hop: 1, NodeId: 52.52.52.52, TeLinkId: 1010702, Loose
LastChange: 02/04/2015 23:30:34
Hop: 2, NodeId: 10.27.255.219, TeLinkId: 1010702, Strict
LastChange: 02/04/2015 23:30:34

Total Path Count: 1
```

## peer

|                    |                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>peer</b> [ <i>peer-node-id</i> ] [ <i>detail</i> ]                                                                                                                |
| <b>Context</b>     | tools>dump>router>gmpls                                                                                                                                              |
| <b>Description</b> | This command dumps information about peer nodes. Configuring the <i>peer-node-id</i> parameter will only dump information about the specified peer node.             |
| <b>Parameters</b>  | <i>peer-node-id</i> — The identifier of the peer node.<br><b>Values</b> a.b.c.d   1 — 4294967295<br><b>detail</b> — Keyword to output information in greater detail. |

### Output

#### Sample Output

```
*A:Dut-B-UNI-C>tools>dump>router>gmpls# peer 10.27.255.213
PEER: vrId 1 PeerAddr 10.27.255.213, AdminState: UP, OperState: UP, HelloState: U
P
Total Peer Count: 1

*A:Dut-B-UNI-C>tools>dump>router>gmpls# peer 10.27.255.213 detail
PEER vrId 1 PeerAddr 10.27.255.213
AdminState: UP, OperState:UP, OperDownReason: N/A
LastChange: 02/04/2015 23:30:34, UpTime: 0d 15:12:39
HelloInterval: 3000msecs, LspHoldTime: 60secs
Flags: 0x30, HelloState: Up
SrcInst: 0xffffffff16, DstInst: 0x54d2ba31, PrevSrcInst: 0xffffffff16
PeerDownCount: 13, InstMismatchCount: 2, TimeoutCount: 0
```

```

Source: 52.52.52.52, Dest: 10.27.255.213 NextHop: 172.21.36.173 [If: 3]
Status: Up, MTU: 0, NumChngInNextHop: 0, LastChange: 02/05/2015 10:53:24

GrHEpoch: 0x5b298, RestartTime: 4294967295, RecoveryTime: 90000
PrevDstInst: 0x0, InvokedCount: 12, ScanEvent: N [N N]

RefreshReduction:
 Local -
MsgId: Enabled, AckDesired: Enabled, Srefresh: Enabled, Epoch: 3236554
 Remote -
MsgId: Supported, Srefresh: Supported, Epoch: 13810201, HighestMsgIdRx:
3670
 NumTxMsgIds: 4, NumRxMsgIds: 4
 NumOutOfOrderMsg: 0, NumRetransmittedMsg: 11, NumPendingAckNack: 0
 NextSrefresh: 9sec

DownstreamSessionCount: 2, UpstreamSessionCount: 2
Path Timeouts: 0, Resv Timeouts: 0

Packet Statistics:
 Hello Tx : 17 Hello Rx : 19338
 Path Tx : 194 Path Rx : 424
 PathErr Tx: 0 PathErr Rx: 183
 PathTear Tx: 5 PathTear Rx: 2
 Resv Tx : 22 Resv Rx : 58
 ResvErr Tx: 2 ResvErr Rx: 0
 ResvTear Tx: 0 ResvTear Rx: 0
 Notify Tx : 0 Notify Rx : 5
 Srfresh Tx: 1946 Srfresh Rx: 1933
 Ack Tx : 2157 Ack Rx : 2170

Total Peer Count: 1

```

# port

|                    |                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>port</b> [ <i>session-name</i> ] [ <i>sender ip-address</i> ] [ <i>detail</i> ]                                                                                                                                                                                                        |
| <b>Context</b>     | tools>dump>router>gmpls                                                                                                                                                                                                                                                                   |
| <b>Description</b> | This command dumps information about a port.                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <i>session-name</i> — The name of the session. 80 characters max; accepts * as a wildcard character.<br><i>sender ip-address</i> — The IP address of the sender, in dotted decimal notation.<br><b>Values</b> a.b.c.d<br><b>detail</b> — Keyword to output information in greater detail. |

## Output

### Sample Output

```
*A:Dut-B-UNI-C>tools>dump>router>gmpls# port sender
```

```

sender <ip-address>
52.52.52.52 42.42.42.42 62.62.62.62 72.72.72.72

*A:Dut-B-UNI-C>tools>dump>router>gmpls# port sender 42.42.42.42

LspPort "vrId: 1 Type: 0 Session: A1::5 Sender: 0.0.0.0"
 TunnelGrpId: 5 MemberId: 5 Mode:LoadSharing
 cfgFarEnd: 42.42.42.42 Flags: 0x0 NumUpdSent: 7 NumSrlgChng: 0
 PSB: vrId: 1 Session (To: 42.42.42.42 - 3 -
52.52.52.52), Sender (52.52.52.52 -
299)
 SRLG Collected :-
 SRLG 27
 SRLG List: Num 2 -->
 4 5
 SRLG 27

LspPort "vrId: 1 Type: 0 Session: B1::1 Sender: 42.42.42.42"
 TunnelGrpId: 1 MemberId: 1 Mode:LoadSharing
 cfgFarEnd: 42.42.42.42 Flags: 0x0 NumUpdSent: 24 NumSrlgChng: 0
 PSB: vrId: 1 Session (To: 52.52.52.52 - 1 -
42.42.42.42), Sender (42.42.42.42 - 46)

LspPort "vrId: 1 Type: 0 Session: B1::4 Sender: 72.72.72.72"
 TunnelGrpId: 4 MemberId: 4 Mode:LoadSharing
 cfgFarEnd: 72.72.72.72 Flags: 0x0 NumUpdSent: 20 NumSrlgChng: 0
 PSB: vrId: 1 Session (To: 52.52.52.52 - 3 - 72.72.72.72), Sender (72.72.72.72
- 286)

LspPort "vrId: 1 Type: 0 Session: B1::6 Sender: 62.62.62.62"
 TunnelGrpId: 6 MemberId: 6 Mode:LoadSharing
 cfgFarEnd: 62.62.62.62 Flags: 0x0 NumUpdSent: 15 NumSrlgChng: 0
 PSB: vrId: 1 Session (To: 52.52.52.52 - 3 - 62.62.62.62), Sender (62.62.62.62
- 871)

LspPort "vrId: 1 Type: 0 Session: C1::3 Sender: 0.0.0.0"
 TunnelGrpId: 3 MemberId: 1 Mode:LoadSharing
 cfgFarEnd: 62.62.62.62 Flags: 0x0 NumUpdSent: 12 NumSrlgChng: 3
 PSB: vrId: 1 Session (To: 62.62.62.62 - 2 - 52.52.52.52), Sender (52.52.52.52
- 297)
 SRLG Collected :-
 SRLG 27
 SRLG 5
 SRLG 27

LspPort "vrId: 1 Type: 0 Session: D1::2 Sender: 0.0.0.0"
 TunnelGrpId: 2 MemberId: 2 Mode:LoadSharing
 cfgFarEnd: 72.72.72.72 Flags: 0x0 NumUpdSent: 10 NumSrlgChng: 0
 PSB: vrId: 1 Session (To: 72.72.72.72 - 1 - 52.52.52.52), Sender (52.52.52.52
- 296)
 SRLG Collected :-
 SRLG 27
 SRLG 4
 SRLG 27

```

```

Total Port Count: 6

*A:Dut-B-UNI-C>tools>dump>router>gmpls# port sender 42.42.42.42
<session-name>
"A1::5" "B1::1" "B1::6" "B1::4" "C1::3" "D1::2"
detail

*A:Dut-B-UNI-C>tools>dump>router>gmpls# port sender 42.42.42.42 detail

LspPort "vrId: 1 Type: 0 Session: A1::5 Sender: 0.0.0.0"
 TunnelGrpId: 5 MemberId: 5 Mode:LoadSharing
 cfgFarEnd: 42.42.42.42 Flags: 0x0 NumUpdSent: 7 NumSrlgChng: 0
 PSB: vrId: 1 Session (To: 42.42.42.42 - 3 - 52.52.52.52), Sender (52.52.52.52
- 299)
 SRLG Collected :-
 SRLG 27
 SRLG List: Num 2 -->
 4 5
 SRLG 27

LspPort "vrId: 1 Type: 0 Session: B1::1 Sender: 42.42.42.42"
 TunnelGrpId: 1 MemberId: 1 Mode:LoadSharing
 cfgFarEnd: 42.42.42.42 Flags: 0x0 NumUpdSent: 24 NumSrlgChng: 0
 PSB: vrId: 1 Session (To: 52.52.52.52 - 1 - 42.42.42.42), Sender (42.42.42.42
- 46)

LspPort "vrId: 1 Type: 0 Session: B1::4 Sender: 72.72.72.72"
 TunnelGrpId: 4 MemberId: 4 Mode:LoadSharing
 cfgFarEnd: 72.72.72.72 Flags: 0x0 NumUpdSent: 20 NumSrlgChng: 0
 PSB: vrId: 1 Session (To: 52.52.52.52 - 3 - 72.72.72.72), Sender (72.72.72.72
- 286)

LspPort "vrId: 1 Type: 0 Session: B1::6 Sender: 62.62.62.62"
 TunnelGrpId: 6 MemberId: 6 Mode:LoadSharing
 cfgFarEnd: 62.62.62.62 Flags: 0x0 NumUpdSent: 15 NumSrlgChng: 0
 PSB: vrId: 1 Session (To: 52.52.52.52 - 3 - 62.62.62.62), Sender (62.62.62.62
- 871)

LspPort "vrId: 1 Type: 0 Session: C1::3 Sender: 0.0.0.0"
 TunnelGrpId: 3 MemberId: 1 Mode:LoadSharing
 cfgFarEnd: 62.62.62.62 Flags: 0x0 NumUpdSent: 12 NumSrlgChng: 3
 PSB: vrId: 1 Session (To: 62.62.62.62 - 2 - 52.52.52.52), Sender (52.52.52.52
- 297)
 SRLG Collected :-
 SRLG 27
 SRLG 5
 SRLG 27

LspPort "vrId: 1 Type: 0 Session: D1::2 Sender: 0.0.0.0"
 TunnelGrpId: 2 MemberId: 2 Mode:LoadSharing
 cfgFarEnd: 72.72.72.72 Flags: 0x0 NumUpdSent: 10 NumSrlgChng: 0
 PSB: vrId: 1 Session (To: 72.72.72.72 - 1 - 52.52.52.52), Sender (52.52.52.52
- 296)

```

```

SRLG Collected :-
 SRLG 27
 SRLG 4
 SRLG 27

```

```

Total Port Count: 6

```

## psb

|                    |                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>psb</b> [ <i>tunnelid tunnel-id</i> ] [ <i>lspid lsp-id</i> ] [ <i>detail</i> ]                                                                                                                                            |
| <b>Context</b>     | tools>dump>router>gmpls                                                                                                                                                                                                       |
| <b>Description</b> | This command dumps information about a PSB.                                                                                                                                                                                   |
| <b>Parameters</b>  | <i>tunnel-id</i> — The identifier of the tunnel.<br><b>Values</b> 0 — 4294967295<br><i>lsp-id</i> — The identifier of the LSP.<br><b>Values</b> 1 — 65535<br><b>detail</b> — Keyword to output information in greater detail. |

## Output

### Sample Output

```

*A:Dut-B-UNI-C>tools>dump>router>gmpls# psb
PSB:
 vrId: 1 Session (To: 42.42.42.42 - 3 - 52.52.52.52), Sender (52.52.52.52 - 299)

PSB:
 vrId: 1 Session (To: 52.52.52.52 - 1 - 42.42.42.42), Sender (42.42.42.42 - 46)

PSB:
 vrId: 1 Session (To: 52.52.52.52 - 3 - 62.62.62.62), Sender (62.62.62.62 - 871)

PSB:
 vrId: 1 Session (To: 52.52.52.52 - 3 - 72.72.72.72), Sender (72.72.72.72 - 286)

PSB:
 vrId: 1 Session (To: 62.62.62.62 - 2 - 52.52.52.52), Sender (52.52.52.52 - 297)

PSB:
 vrId: 1 Session (To: 72.72.72.72 - 1 - 52.52.52.52), Sender (52.52.52.52 - 296)

Total PSB Count: 6

*A:Dut-B-UNI-C>tools>dump>router>gmpls# psb tunnelid 1 detail

PSB:
 vrId: 1 Session (To: 52.52.52.52 - 1 - 42.42.42.42), Sender (42.42.42.42 - 46)

PSB CurrState: GMPLS_PSB_UP PrevState: GMPLS_PSB_UP Flags: 0x0

```

```

isIngress: N isTransit: N isEgress: Y
DnStream Peer: None
UpStream Peer: vrId 1 PeerAddr 10.27.255.213
UpStream TELink: vrId 1 TELinkId 1010701, dBLink: vrId 1 DBLinkId 101070117
 RemoteDBLinkId: 101070117, PortId: 0x2268000, refCnt: 1, State: UP

Sender Template - Sender: 42.42.42.42, LspId: 46
Session Attribute -
 Session Name: B1::1
 CType: 7, HoldPri: 1, SetupPri: 5, Flags: 0x2
 IncludeGroup: 0x0, IncludeAllGroup: 0x0, ExcludeGroup: 0x0
Lsp Attribute - TLV: 1, Flags: 0x80000
NextPathRefresh: 0secs, PathRefreshTimeout: 151secs
Path RX Message Id: 2209, Epoch: 13810201, Flags: 0x1

PSB RRO : ->
(1) Unnumbered RouterId 10.27.255.213, InterfaceId 1010701, Flags 0x0
(2) SRLG 27
(3) Label 101070117, Flags 0x0
(4) Unnumbered RouterId 10.27.255.213, InterfaceId 13004, Flags 0x0
(5) Label 101070130, Flags 0x80
(6) Unnumbered RouterId 10.27.255.215, InterfaceId 13007, Flags 0x0
(7) SRLG 6
(8) Label 101070130, Flags 0x0
(9) Unnumbered RouterId 10.27.255.215, InterfaceId 1010701, Flags 0x0
(10) SRLG 27
(11) Label 101070117, Flags 0x80
(12) Unnumbered RouterId 42.42.42.42, InterfaceId 1010701, Flags 0x0
(13) Label 101070117, Flags 0x0

PSB ERO : ->
(1) IPv4Prefix 52.52.52.52/32, Strict [1]

PSB XRO : ->
 NULL XRO
PROTECTION - Flags: 0x0, E2EProtectionType: 0

 LinkFlags: 0x0, SegProtFlags: 0x0, SegProtType: 0x1
ASSOCIATION - Type: 0, AssociationId: 46, Source: 42.42.42.42
NOTIFY REQ - Node: 10.27.255.213

RSB:
NextResvRefresh: 9secs, ResvRefreshTimeout: 0secs
Resv TX Message Id: 2661, Flags: 0x0
PROTECTION - Flags: 0x0, E2EProtectionType: 0
 LinkFlags: 0x0, SegProtFlags: 0x0, SegProtType: 0x1

FLOWSPEC : ->
 CType: 6 SwitchingGranularity: 1, MTU: 9212, TlvFlags: 0x1
 EthBWProfileTlv - Profile: 0, Index: 0
 CIR: 10.000 Gbps, CBS: 10.000 Gbps
 EIR: infinity, EBS: infinity
 FILTERSPEC Label : 0
 FILTERSPEC RRO : ->
(1) Unnumbered RouterId 52.52.52.52, InterfaceId 1010701, Flags 0x0
(2) Label 101070117, Flags 0x80

NOTIFY REQ - Node: 52.52.52.52

```



```

Bound to LSPPORT: vrId: 1 Type: 0 Session: B1::1 Sender: 42.42.42.42

Num Paths Received :78
Num Paths Transmitted :0
Num Resvs Received :0
Num Resvs Transmitted :1

Num Summmary Paths Received :737
Num Summmary Paths Transmitted:0
Num Summmary Resvs Received :0
Num Summmary Resvs Transmitted:750
Created at 37048 (21983 secs back)

PSB:
 vrId: 1 Session (To: 72.72.72.72 - 1 - 52.52.52.52), Sender (52.52.52.52 - 296)

PSB CurrState: GMPLS_PSB_UP PrevState: GMPLS_PSB_UP Flags: 0x0
isIngress: Y isTransit: N isEgress: N
DnStream Peer: vrId 1 PeerAddr 10.27.255.213
UpStream Peer: None
DnStream TElLink: vrId 1 TElLinkId 1010702, dBLink: vrId 1 DBLinkId 101070217
 RemoteDBLinkId: 101070217, PortId: 0x2270000, refCnt: 1, State: UP

Sender Template - Sender: 52.52.52.52, LspId: 296
Session Attribute -
 Session Name: D1::2
 CType: 7, HoldPri: 1, SetupPri: 5, Flags: 0x2
 IncludeGroup: 0x0, IncludeAllGroup: 0x0, ExcludeGroup: 0x0
Lsp Attribute - TLV: 1, Flags: 0x80000
NextPathRefresh: 3secs, PathRefreshTimeout: 0secs
Path TX Message Id: 471, Flags: 0x0

PSB RRO : ->
 (1) Unnumbered RouterId 52.52.52.52, InterfaceId 1010702, Flags 0x0
 (2) Label 101070217, Flags 0x0

PSB ERO : ->
 (1) Unnumbered RouterId 52.52.52.52, LinkId 1010702, Loose [1]
 (2) IPv4Prefix 10.27.255.213/32, Strict [0]
 (3) Unnumbered RouterId 10.27.255.219, LinkId 1010702, Strict [0]
 (4) IPv4Prefix 72.72.72.72/32, Loose [0]

PSB XRO : ->
 NULL XRO
PROTECTION - Flags: 0x0, E2EProtectionType: 1
 LinkFlags: 0x0, SegProtFlags: 0x0, SegProtType: 0x0
ASSOCIATION - Type: 0, AssociationId: 296, Source: 52.52.52.52
NOTIFY REQ - Node: 52.52.52.52

RSB:
NextResvRefresh: 0secs, ResvRefreshTimeout: 151secs
Resv RX Message Id: 19, Epoch: 13810201, Flags: 0x1
PROTECTION - Flags: 0x0, E2EProtectionType: 1
 LinkFlags: 0x0, SegProtFlags: 0x0, SegProtType: 0x0

FLOWSPEC : ->
 CType: 6 SwitchingGranularity: 1, MTU: 9212, TlvFlags: 0x1
 EthBWProfileTlv - Profile: 0, Index: 0

```

```

CIR: 10.000 Gbps, CBS: 10.000 Gbps
EIR: infinity, EBS: infinity
FILTERSPEC Label : 101070217
FILTERSPEC RRO : ->
(1) Unnumbered RouterId 52.52.52.52, InterfaceId 1010702, Flags 0x0
(2) Label 101070217, Flags 0x0
(3) Unnumbered RouterId 10.27.255.213, InterfaceId 1010702, Flags 0x0
(4) SRLG 27
(5) Label 101070217, Flags 0x80
(6) Unnumbered RouterId 10.27.255.213, InterfaceId 13000, Flags 0x0
(7) SRLG 4
(8) Label 101070230, Flags 0x0
(9) Unnumbered RouterId 10.27.255.219, InterfaceId 13001, Flags 0x0
(10) Label 101070230, Flags 0x80
(11) Unnumbered RouterId 10.27.255.219, InterfaceId 1010702, Flags 0x0
(12) SRLG 27
(13) Label 101070217, Flags 0x0
(14) Unnumbered RouterId 72.72.72.72, InterfaceId 1010702, Flags 0x0
(15) Label 101070217, Flags 0x80

NOTIFY REQ - Node: 10.27.255.213

Bound to LSPPORT: vrId: 1 Type: 0 Session: D1::2 Sender: 0.0.0.0

Num Paths Received :0
Num Paths Transmitted :1
 Num Resvs Received :23
Num Resvs Transmitted :0

Num Summmary Paths Received :0
Num Summmary Paths Transmitted:1834
Num Summmary Resvs Received :1824
Num Summmary Resvs Transmitted:0
Created at 4439 (54592 secs back)

Total PSB Count: 2
```

status

|                    |                                                      |
|--------------------|------------------------------------------------------|
| <b>Syntax</b>      | <b>status</b>                                        |
| <b>Context</b>     | tools>dump>router>gmpls                              |
| <b>Description</b> | This command dumps general GMPLS status information. |
| <b>Output</b>      |                                                      |

Sample Output

```
*A:Dut-B-UNI-C>tools>dump>router>gmpls# status

GMPLS instance vrId: 1
AdminState: UP, OperState:UP, OperDownReason: N/A
Flags: 0x0, localNodeId: 52.52.52.52
InitRetryTimeout: 30 KeepMultiplier: 3 RefreshTime: 30
```

```
RapidRetransmitTime: 5 RapidRetryLimit: 3
GrRestartTime: 180 GrRecoveryTime: 300

gLspWPOriginate: 3 gLspWPTransit: 0 gLspWPTerminate: 3
gLspPPOriginate: 0 gLspPPTransit: 0 gLspPPTerminate: 0
NumTELink: 6 NumDB: 6 NumLspPort: 6 NumTunGrp: 6
Num gLsp: 3 Num gLspPath: 3 Num Path: 4
NumLspPortAudited: 0 NumStaleLspPortDeleted: 0

```

## te-link

|                    |                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>te-link</b> [ <i>te-link-id</i> ] [ <i>detail</i> ]                                                                                                                       |
| <b>Context</b>     | tools>dump>router>gmpls                                                                                                                                                      |
| <b>Description</b> | This command dumps information about TE links. Configuring the <i>te-link-id</i> parameter will only dump information about the specified TE link.                           |
| <b>Parameters</b>  | <i>te-link-id</i> — The identifier of the TE link.<br><b>Values</b> 1 — 4294967295   <i>te-link-name</i><br><b>detail</b> — Keyword to output information in greater detail. |

### Output

#### Sample Output

```
*A:Dut-B-UNI-C>tools>dump>router>gmpls# te-link
<te-link-id>
1010701 1010702 1010703 1010704 1010705 1010706
detail

*A:Dut-B-UNI-C>tools>dump>router>gmpls# te-link 1010701

TE-LINK "vrId 1 TELinkId 1010701"
AdminState: UP, OperState:UP
LastChange: 02/04/2015 23:30:34
Peer: 10.27.255.213 RemoteTELinkId: 1010701 NumDbLink: 1

Total TE-Link Count: 1

*A:Dut-B-UNI-C>tools>dump>router>gmpls# te-link 1010701 detail

TE-LINK "vrId 1 TELinkId 1010701"
AdminState: UP, OperState:UP
LastChange: 02/04/2015 23:30:34
Peer: 10.27.255.213 RemoteTELinkId: 1010701 NumDbLink: 1

Total TE-Link Count: 1
```



---

## 4 PCEP

### 4.1 In This Chapter

This chapter provides information to configure MPLS and RSVP.

- [Introduction to the Path Computation Element Protocol \(PCEP\)](#)
  - [PCC and PCE Configuration](#)
  - [Base Implementation of Path Computation Elements \(PCE\)](#)
  - [PCEP Session Establishment and Maintenance](#)
  - [PCEP Parameters](#)
- [PCEP Configuration Command Reference](#)

### 4.2 Introduction to the Path Computation Element Protocol (PCEP)

The Path Computation Element Protocol (PCEP) is one of several protocols used for communication between a Wide-Area Network (WAN) Software-Define Networking (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 implements a common set of standards-based south-bound interfaces to the network elements for both topology discovery and tunnel and flow programming. It is a virtual SR OS (vSROS) image which applies the south-bound interfaces to the network elements and the adaptation layer to the applications. The south-bound interfaces include IGP and BGP-LS 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 which can be used together or separately as illustrated in [Figure 53](#).

**Figure 53 NSP Functional Modules**



0974


The two main features of the NSP are as follows:

- **Network Services Director (NSD)** — The NSD is a programmable and multi-vendor service provisioning tool exposing a single and simple API to the user and OSS. It implements service model abstraction and adapts to each vendor-specific service model. It supports provisioning services such as ELINE, ELAN, ETREE, L3 VPN, traffic steering, and service chaining.

- **Network Resource Controller (NRC)** — The NRC implements a separate module 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 NRC-P implements the stateful Path Computation Element (PCE) for packet networks. [Figure 54](#) illustrates the NRC-P architecture and its main components.

**Figure 54 Packet Network Resource Controller (NRC-P) Architecture**



0973

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 handles the functions of topology discovery of multiple IGP instances and areas via IGP or BGP-LS and the PCE PCEP functions
- 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, such as the PCE Client (PCC). It also uses the 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 or RSVP LSPs, the PCE owns the path updating and periodic re-optimization of the LSP. In this case, the PCE acts in an active stateful role. The PCE can also act in a stateful passive role for other LSPs on the router by discovering them 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, per RFC 5440
- active and passive stateful PCE LSP update, as per draft-ietf-pce-stateful-pce
- delegation of LSP control to PCE
- synchronization of the LSP database with network elements for PCE-controlled LSPs and network element-controlled LSPs
- support for the SR-TE P2P LSP type, as per draft-ietf-pce-segment-routing
- 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 bi-directionality constraints using extensions to the PCE path profile, as per draft-alvarez-pce-path-profiles

## 4.2.1 PCC and PCE Configuration

The following PCE parameters cannot be modified while the PCEP session is operational:

- **local-address**
- **keepalive**
- **dead-timer**



The **unknown-message-rate** PCE parameter can be modified while the PCEP session is operational.

The following PCC parameters cannot be modified while the PCEP session is operational:

- **local-address**
- **keepalive**
- **dead-timer**
- **peer** (regardless of **shutdown** state)

The following PCC parameters can be modified while the PCEP session is operational:

- **report-path-constraints**
- **unknown-message-rate**

## 4.2.2 Base Implementation of Path Computation Elements (PCE)

The base implementation of PCE uses the PCEP extensions defined in RFC 5440 [pce-pcep].

The main functions of the 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 PCC or path computation cancellation by PCE
- error messages (PCErr) by which the PCEP speaker can inform its peer about errors related to processing requests, message objects, or TLVs

[Table 34](#) lists the messages and objects that are supported in SR OS Release 14.0.

**Table 34 Base PCEP Message Objects and TLVs**

| TLV, Object, or Message | Contained in Object | Contained in Message |
|-------------------------|---------------------|----------------------|
| OPEN Object             | N/A                 | OPEN, PCErr          |

**Table 34 Base PCEP Message Objects and TLVs (Continued)**

| TLV, Object, or Message       | Contained in Object | Contained in Message             |
|-------------------------------|---------------------|----------------------------------|
| 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 <sup>1</sup> |
| METRIC Object                 | N/A                 | PCReq, PCRep, PCRpt <sup>1</sup> |
| Explicit Route Object (ERO)   | N/A                 | PCRep                            |
| Reported Route Object (RRO)   | N/A                 | PCReq                            |
| LSPA Object                   | N/A                 | PCReq, PCRep, PCRpt <sup>1</sup> |
| 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 behavior and limitations of the implementation of the objects in [Table 34](#) are as follows:

- 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. PCE returns the computed value in PCRep with flags set identically to PCReq.
- 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. PCE returns the computed value in PCRep with flags set identically to PCReq.

- The IRO and LOAD-BALANCING objects are not in 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.
- LSP path setup and hold priorities will be configurable during SR-TE LSP configuration on the router, and PCC will pass the configurations on in an LSPA object. However, PCE does not implement LSP pre-emption.
- 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 in the SR OS:

- 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 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 PCE. This is not specific to PCE controlling an SR-TE LSP, but also to controlling an RSVP LSP.
- 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

### 4.2.3 PCEP Session Establishment and Maintenance

The PCEP protocol operates over TCP using destination TCP port 4189. The PCC always initiates the connection. Once the user configures the PCEP local address and the peer address on the PCC, the PCC initiates a TCP connection to the PCE. Once a connection is established, the PCC and PCE exchange OPEN messages, which initializes the PCEP session and exchanges the session parameters to be negotiated.


The PCC always checks first if the remote PCE address is reachable out-of-band via the management port. If not, it will try to reach the remote PCE address in-band. When the session comes up out-of-band, the system IP address is always used; the local address configured by the user is ignored and is only used for an in-band session.

A keep-alive mechanism is used as an acknowledgment of the acceptance of the session within the negotiated parameters. It is also used as a maintenance function to detect whether or not the PCEP peer is still alive.

The negotiated parameters include the Keepalive timer and the DeadTimer, and one or more PCEP capabilities such as support of Stateful PCE and the SR-TE LSP Path type.

The PCEP session initialization steps are illustrated in [Figure 55](#).

**Figure 55 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.

## 4.2.4 PCEP Parameters

The following PCEP parameters are user-configurable on both the PCC and PCE. On the PCE, the configured parameter values are used on sessions to all PCCs.

- **Keep-alive timer** — A PCEP speaker (PCC or PCE) must send a keep-alive 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 keep-alive message is sent.

The keep-alive mechanism is asymmetric, meaning that each peer can use a different keep-alive timer value.

The range of this parameter is 1 to 255 seconds, and the default value is 30 seconds. The no version returns to the default value.

- **Dead timer** — This timer tracks the amount of time a PCEP speaker (PCC or PCE) 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 parameter is 1 to 255 seconds, and the default value is 120 seconds. The no version returns to the default value.

- **Maximum rate of unknown messages** — When the rate of received unrecognized or unknown messages reaches this limit, the PCEP speaker closes the session to the peer.
- **Session re-delegation 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 will continue to try re-establishing the PCEP session. When the PCEP session is re-established, the LSP database is synchronized with the PCE, and any LSP which went down since the last time the PCEP session was up will have its path updated by the PCE.

### 4.2.4.1 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 re-optimization and updating of the LSP paths.

[Table 35](#) describes the messages and objects supported by stateful PCE in the SR OS.

**Table 35 PCEP Stateful PCE Extension Objects and TLVs**

| 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 <sup>1</sup> , PCErr        |
| LSP Object                              | ERO                 | PCRpt <sup>1</sup> , PCReq, PCRep |
| LSP Identifiers TLV                     | LSP                 | PCRpt <sup>1</sup>                |
| Symbolic Path Name TLV                  | LSP, SRP            | PCRpt <sup>1</sup>                |
| LSP Error Code TLV                      | LSP                 | PCRpt <sup>1</sup>                |
| RSVP Error Spec TLV                     | LSP                 | PCRpt <sup>1</sup>                |

Note:

1. Nokia proprietary

The behavior and limitations of the implementation of the objects in [Table 35](#) are as follows:

- 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 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 which are signaled by the routers but are not synchronized up 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 PCE will be able to bring the PCEP session up, the LSP database will not be updated, since stateless PCC does not support the PCRpt message. As such, PCE path computation will not accurately take into account the bandwidth used by these LSPs in the network.
  - PCE ignores the R-flag (re-optimize flag) in the PCReq message when acting in stateful-passive mode for a given LSP, and will always return the new computed path, regardless if it is link-by-link identical or has the same metric as the current path. The decision whether or not to initiate the new path in the network belongs to the PCC.
  - The SVEC object is not supported in the SR OS 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 PCE to become aware of the original constraints of such an LSP, the following additional procedures are performed. These procedures are proprietary to Nokia.
    - 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. SR OS PCC does not support putting a bound 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 contains a single set. If the PCRpt contains a duplicate set, PCE path computation must use the constraints in the duplicate set.
    - For interoperability, implementations compliant to PCEP standards should be able to accept the first metric object and ignore the second object without additional error handling. Since 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.

In a future release, when the stateful PCE supports LSP initiation, the extensions described in [Table 36](#) are required.

**Table 36      PCEP PCE-Initiated LSP Extension Objects and TLVs**

| TLV, Object, or Message               | Contained in Object | Contained in Message |
|---------------------------------------|---------------------|----------------------|
| PCE LSP Initiate Message (PCInitiate) | N/A                 | New message          |
| PCC LSP Create Flag (C-flag)          | LSP                 | PCRpt <sup>1</sup>   |

Note:

- 1. Nokia proprietary

**4.2.4.2      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.

- A new 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.
- 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.

- 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 as 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 Segment-ID (SID) Depth value in a METRIC object for both a PCE-computed LSP and a PCE-controlled LSP. 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 Segment-ID (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's 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.

Table 37 describes the segment routing extension objects and TLVs supported in the SR OS.

**Table 37 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 <sup>1</sup> |
| SR-ERO Sub-object                             | ERO                 | PCRep, PCRpt <sup>1</sup>        |
| SR-RRO Sub-object                             | RRO                 | PCReq, PCRpt <sup>1</sup>        |
| Segment-ID (SID) Depth Value in METRIC Object | METRIC              | PCReq, PCRpt <sup>1</sup>        |

Note:

1. Nokia proprietary

#### 4.2.4.3 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 SR OS support three different modes of operations for PCC-initiated LSPs which are configurable on a per-LSP basis.

- When the path of the LSP is computed and updated by the router acting as a PCE Client (PCC), the LSP is referred to as PCC-initiated and PCC-controlled. A PCC-initiated and PCC-controlled LSP has the following characteristics:
  - The LSP can contain strict or loose hops, or a combination of both.
  - CSPF is supported for RSVP-TE LSPs. Local path computation takes the form of hop-to-label translation for SR-TE 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 can not be updated by the PCE. In other words, 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 LSP has the following characteristics:

- The user must enable the **pce-computation** option for the LSP so that the PCE can perform path computation at the request of the PCC only. 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.
- 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 user must enable the **pce-control** option for the LSP so that the PCE can perform path updates following a network event without an explicit request from the PCC. PCC delegates full control.
- The user must enable the **pce-report** option for LSPs that cannot be delegated to the PCE. The PCE acts in active stateful mode for this LSP.

#### 4.2.4.3.1 PCC-Initiated and PCE-Computed/Controlled LSPs

The following is the procedure for configuring and programming a PCC-initiated SR-TE LSP when control is delegated to the PCE.

- Step 1.** The LSP configuration is created on the PE router via CLI or via the OSS/5620 SAM.
- The configuration dictates which PCE control mode is desired: active (**pce-control** and **pce-report** options enabled) or passive (**pce-computation** enabled and **pce-control** disabled).
- Step 2.** PCC assigns a unique PLSP-ID to the LSP. The PLSP-ID uniquely identifies the LSP on a PCEP session and must remain constant during its lifetime. PCC on the router must keep track of the association of the PLSP-ID to the Tunnel-ID and Path-ID, and use the latter to communicate with MPLS about a specific path of the LSP. PCC also uses the SRP-ID to correlate PCRpt messages for each new path of the LSP.

- Step 3.** The PE router does not validate the entered path. Note however that in the SR OS, 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.
- Step 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.
- Step 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.

- Step 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.
- Step 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.
- Step 8.** Following a network event or a re-optimization, the PCE computes a new path for a PCE-controlled LSP and returns it in a PCUpd message with the new ERO. It will include the LSP object with the same unique PLSP-ID assigned by the PCC, as well as the Stateful Request Parameter (SRP) object with a unique SRP-ID-number to track error and state messages specific to this new path.
- Step 9.** The PE router updates the CPM and the data path with the new path.
- Step 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.

- Step 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.
- Step 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.
- Step 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:

- a. If the user enabled the **pce-computation** option on a PCC-controlled LSP with an active path, no action is performed until the next time the router needs a path for the LSP following a network event of a LSP parameter change. At that point, the prior procedure is followed.
- b. 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 re-optimization. At that point, the prior procedure is followed.

#### 4.2.4.3.2 LSP Path Diversity and Bi-Directionality Constraints

The PCE path profile defined in draft-alvarez-pce-path-profiles is used to request path diversity or a disjoint for two or more LSPs originating on the same or different PE routers. It is also used to request that paths of two unidirectional LSPs between the same two routers use the same TE links. This is referred to as the bi-directionality constraint.

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 bound 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 bi-directionality, symmetric reverse route preferred, symmetric reverse route required
- maximum path IGP metric (cost)
- maximum path TE metric
- maximum hop count

In addition, the user can specify in the PCE Path Profile which PCE objective to use to optimize the path of the LSP.

- IGP metric (cost)
- TE metric
- Hops (span)

The CSPF algorithm will optimize this objective. If a constraint is provided for the same metric, then the CSPF algorithm makes sure that the selected path achieves a lower or equal value to the bound 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/exclude admin-group constraints and setup and hold priorities, are not supported in the PCE profile.

In order to indicate the path diversity and bi-directionality 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 for this is described in the [Configuring and Operating SR-TE](#) section. 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 which 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 a given 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 into 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 PCC to PCE.

The bi-directionality constraint operates the same way as the diversity constraint. The user can configure a PCE profile with both the path diversity and bi-directionality constraints. PCE will check if there is an LSP in the reverse direction which belongs to the same path group ID as an originating LSP it is computing the path for, and will enforce the constraint.

In order to for PCE to be aware of the path diversity and bi-directionality constraints for an LSP which 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.

Table 38 describes the new objects introduced in the PCE path profile.

**Table 38      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 <sup>1</sup> |

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.

## 4.3 PCEP Configuration Command Reference

### 4.3.1 Command Hierarchies

- [PCEP Commands](#)

#### 4.3.1.1 PCEP Commands

```
config
 — router
 — [no] pcep
 — pcc
 — dead-timer seconds
 — no dead-timer
 — keepalive seconds
 — no keepalive
 — local-address ip-address
 — no local-address
 — [no] peer ip-address
 — [no] report-path-constraints
 — [no] shutdown
 — unknown-message-rate integer
 — no unknown-message-rate
 — pce
 — dead-timer seconds
 — no dead-timer
 — keepalive seconds
 — no keepalive
 — local-address ip-address
 — no local-address
 — [no] shutdown
 — unknown-message-rate integer
 — no unknown-message-rate
```

---

## 4.3.2 Command Descriptions

- [PCEP Commands](#)

### 4.3.2.1 PCEP Commands

#### pcep

|                    |                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] pcep</b>                                                                                                                                                                           |
| <b>Context</b>     | config>router                                                                                                                                                                              |
| <b>Description</b> | This command enables Path Computation Element communications Protocol (PCEP), and enters the context to configure PCEP parameters.<br><br>The <b>no</b> form of the command disables PCEP. |

#### pcc

|                    |                                                               |
|--------------------|---------------------------------------------------------------|
| <b>Syntax</b>      | <b>pcc</b>                                                    |
| <b>Context</b>     | config>router>pcep                                            |
| <b>Description</b> | This command enables the context to configure PCC parameters. |

#### pce

|                    |                                                               |
|--------------------|---------------------------------------------------------------|
| <b>Syntax</b>      | <b>pce</b>                                                    |
| <b>Context</b>     | config>router>pcep                                            |
| <b>Description</b> | This command enables the context to configure PCE parameters. |

#### dead-timer

|                |                                                          |
|----------------|----------------------------------------------------------|
| <b>Syntax</b>  | <b>dead-timer</b> <i>seconds</i><br><b>no dead-timer</b> |
| <b>Context</b> | config>router>pcep>pcc<br>config>router>pcep>pce         |



---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |               |          |                |     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|----------|----------------|-----|
| <b>Description</b> | <p>This command configures the PCEP session dead timer value, which is the amount of time a PCEP speaker (PCC or PCE) will wait after the receipt of the last PCEP message before declaring its peer down.</p> <p>The keep-alive mechanism is asymmetric, meaning that each PCEP speaker can propose a different dead timer value to its peer to use to detect session timeout.</p> <p>The <b>no</b> form of the command returns the dead timer to the default value.</p> |               |          |                |     |
| <b>Parameters</b>  | <p><i>seconds</i> — the dead timer value, in seconds</p> <table> <tr> <td><b>Values</b></td><td>1 to 255</td></tr> <tr> <td><b>Default</b></td><td>120</td></tr> </table>                                                                                                                                                                                                                                                                                                 | <b>Values</b> | 1 to 255 | <b>Default</b> | 120 |
| <b>Values</b>      | 1 to 255                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |               |          |                |     |
| <b>Default</b>     | 120                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |               |          |                |     |

## keepalive

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |               |          |                |    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|----------|----------------|----|
| <b>Syntax</b>      | <b>keepalive</b> <i>seconds</i><br><b>no keepalive</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                           |               |          |                |    |
| <b>Context</b>     | config>router>pcep>pcc<br>config>router>pcep>pce                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |               |          |                |    |
| <b>Description</b> | <p>This command configures the PCEP session keep-alive value. A PCEP speaker (PCC or PCE) must send a keep-alive 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 keep-alive message is sent.</p> <p>The keep-alive mechanism is asymmetric, meaning that each peer can use a different keep-alive timer value at its end.</p> <p>The <b>no</b> form of the command returns the keep-alive timer to the default value.</p> |               |          |                |    |
| <b>Parameters</b>  | <p><i>seconds</i> — the keep-alive value, in seconds</p> <table> <tr> <td><b>Values</b></td><td>1 to 255</td></tr> <tr> <td><b>Default</b></td><td>30</td></tr> </table>                                                                                                                                                                                                                                                                                                                                         | <b>Values</b> | 1 to 255 | <b>Default</b> | 30 |
| <b>Values</b>      | 1 to 255                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |               |          |                |    |
| <b>Default</b>     | 30                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |               |          |                |    |

## local-address

|                    |                                                                       |
|--------------------|-----------------------------------------------------------------------|
| <b>Syntax</b>      | <b>local-address</b> <i>ip-address</i><br><b>no local-address</b>     |
| <b>Context</b>     | config>router>pcep>pcc<br>config>router>pcep>pce                      |
| <b>Description</b> | <p>This command configures the local address of the PCEP speaker.</p> |

The PCEP protocol operates over TCP using destination TCP port 4189. The PCE client (PCC) always initiates the connection. Once the user configures the PCEP local address and the peer address on the PCC, the latter initiates a TCP connection to the PCE. When the connection is established, the PCC and PCE exchange OPEN messages, which initializes the PCEP session and exchanges the session parameters to be negotiated.

The PCC always checks first if the remote PCE address is reachable out-of-band via the management port. If not, it will check if the remote PCE address is reachable in-band. When the session comes up out-of-band, the system IP address is always used. The local address configured by the user is only used for in-band sessions, and is otherwise ignored.

The **no** form of 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

## peer

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

**Context** config>router>pcep>pcc  
config>router>pcep>pce

**Description** This command configures the IP address of a peer PCEP speaker. It is used as the destination address in the PCEP session messages to a PCEP peer.

The **no** form of the command removes the specified peer PCEP speaker.

**Parameters** *ip-address* — the IP address of the PCEP peer to be used as the destination address in the PCEP session

## report-path-constraints

**Syntax** **report-path-constraints**  
**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 which is delegated, but for which there is no prior state in its LSP database, such as if no PCReq message was sent for the same PLSP-ID, the following proprietary behavior is observed:

- 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 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 SR OS 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 contained a single set. If the PCRpt 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>pce

**Description** This command administratively disables the PCC or PCE process.

The following PCE parameters can only be modified when the PCEP session is shut down:

- **local-address**
- **keepalive**
- **dead-timer**

The **unknown-message-rate** PCE parameter can be modified without shutting down the PCEP session.

The following PCC parameters can only be modified when the PCEP session is shut down:

- **local-address**
- **keepalive**
- **dead-timer**
- **peer**

The following PCC parameters can be modified without shutting down the PCEP session:

- **report-path-constraints**
- **unknown-message-rate**

**Default** shutdown

---

## unknown-message-rate

|                    |                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>unknown-message-rate</b> <i>integer</i><br><b>no unknown-message-rate</b>                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>pcep>pcc<br>config>router>pcep>pce                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command configures the maximum rate of unknown messages which can be received on a PCEP session.</p> <p>When the rate of received unrecognized or unknown messages reaches the configured limit, the PCEP speaker closes the session to the peer.</p> <p>The <b>no</b> form of the command returns the unknown message rate to the default value.</p> |
| <b>Parameters</b>  | <i>integer</i> — the rate of unknown messages, in messages per minute                                                                                                                                                                                                                                                                                         |
|                    | <b>Values</b> 1 to 255                                                                                                                                                                                                                                                                                                                                        |
|                    | <b>Default</b> 10                                                                                                                                                                                                                                                                                                                                             |

## 5 Label Distribution Protocol

### 5.1 In This Chapter

This chapter provides information to enable Label Distribution Protocol (LDP):

- [Label Distribution Protocol](#)
- [TTL Security for BGP and LDP](#)
- [ECMP Support for LDP](#)
- [Unnumbered Interface Support in LDP](#)
- [LDP over RSVP Tunnels](#)
- [LDP over RSVP Without Area Boundary](#)
- [Class-Based Forwarding of LDP Prefix Packets over IGP Shortcuts](#)
- [LDP ECMP Uniform Failover](#)
- [LDP Fast-Reroute for IS-IS and OSPF Prefixes](#)
- [mLDP Fast Upstream Switchover](#)
- [LDP FEC to BGP Label Route Stitching](#)
- [LDP-SR Stitching for IPv4 prefixes \(IS-IS\)](#)
- [LDP FRR Remote LFA Backup using SR Tunnel for IPv4 Prefixes \(IS-IS\)](#)
- [Automatic Creation of a Targeted Hello Adjacency and LDP Session](#)
- [Multicast P2MP LDP for GRT](#)
- [LDP P2MP Support](#)
- [Multicast LDP Fast Upstream Switchover](#)
- [Multi-Area and Multi-Instance Extensions to LDP](#)
- [LDP Graceful Handling of Resource Exhaustion](#)
- [LDP Enhanced Graceful Handling of Resources](#)
- [User Guidelines and Troubleshooting Procedures](#)
- [LDP IPv6 Control and Data Planes](#)
- [LDP Process Overview](#)
- [LDP-IGP Synchronization](#)

## 5.2 Label Distribution Protocol

Label Distribution Protocol (LDP) is a protocol used to distribute labels in non-traffic-engineered applications. LDP allows routers to establish label switched paths (LSPs) through a network by mapping network-layer routing information directly to data link layer-switched paths.

An LSP is defined by the set of labels from the ingress Label Switching Router (LSR) to the egress LSR. LDP associates a Forwarding Equivalence Class (FEC) with each LSP it creates. A FEC is a collection of common actions associated with a class of packets. When an LSR assigns a label to a FEC, it must let other LSRs in the path know about the label. LDP helps to establish the LSP by providing a set of procedures that LSRs can use to distribute labels.

The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each LSR splices incoming labels for a FEC to the outgoing label assigned to the next hop for the given FEC. The next-hop for a FEC prefix is resolved in the routing table. LDP can only resolve FECs for IGP and static prefixes. LDP does not support resolving FECs of a BGP prefix.

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).

SR OS supports IPv4 and IPv6 in LDP control and data planes; as IPv6 has been added subsequently, CLI commands have changed to support both IPv4 and IPv6. Refer to the *Release 13.0.R1 SR OS Software Release Notes* for more information.

### 5.2.1 LDP and MPLS

LDP performs the label distribution only in MPLS environments. The LDP operation begins with a hello discovery process to find LDP peers in the network. LDP peers are two LSRs that use LDP to exchange label/FEC mapping information. An LDP session is created between LDP peers. A single LDP session allows each peer to learn the other's label mappings (LDP is bi-directional) and to exchange label binding information.

LDP signaling works with the MPLS label manager to manage the relationships between labels and the corresponding FEC. For service-based FECs, LDP works in tandem with the Service Manager to identify the virtual leased lines (VLLs) and Virtual Private LAN Services (VPLSs) to signal.

An MPLS label identifies a set of actions that the forwarding plane performs on an incoming packet before discarding it. The FEC is identified through the signaling protocol (in this case, LDP) and allocated a label. The mapping between the label and the FEC is communicated to the forwarding plane. In order for this processing on the packet to occur at high speeds, optimized tables are maintained in the forwarding plane that enable fast access and packet identification.

When an unlabeled packet ingresses the router, classification policies associate it with a FEC. The appropriate label is imposed on the packet, and the packet is forwarded. Other actions that can take place before a packet is forwarded are imposing additional labels, other encapsulations, learning actions, etc. When all actions associated with the packet are completed, the packet is forwarded.

When a labeled packet ingresses the router, the label or stack of labels indicates the set of actions associated with the FEC for that label or label stack. The actions are preformed on the packet and then the packet is forwarded.

The LDP implementation provides DOD, DU, ordered control, liberal label retention mode support.

## 5.2.2 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 56](#). This illustration shows the interaction of the LDP subsystem with other subsystems, including memory management, label management, service management, SNMP, interface management, and RTM. In addition, debugging capabilities are provided through the logger.


Communication within LDP tasks is typically done by inter-process communication through the event queue, as well as through updates to the various data structures. The primary data structures that LDP maintains are:

- FEC/label database — Contains all FEC to label mappings that include both sent and received. It also contains both address FECs (prefixes and host addresses) and service FECs (L2 VLLs and VPLS)
- Timer database — Contains all timers for maintaining sessions and adjacencies
- Session database — Contains all session and adjacency records, and serves as a repository for the LDP MIB objects

### 5.2.3 Subsystem Interrelationships

The sections below describe how LDP and the other subsystems work to provide services. [Figure 56](#) shows the interrelationships among the subsystems.

**Figure 56 Subsystem Interrelationships**





### **5.2.3.1 Memory Manager and LDP**

LDP does not use any memory until it is instantiated. It pre-allocates some amount of fixed memory so that initial startup actions can be performed. Memory allocation for LDP comes out of a pool reserved for LDP that can grow dynamically as needed. Fragmentation is minimized by allocating memory in larger chunks and managing the memory internally to LDP. When LDP is shut down, it releases all memory allocated to it.

### **5.2.3.2 Label Manager**

LDP assumes that the label manager is up and running. LDP will abort initialization if the label manager is not running. The label manager is initialized at system boot-up; hence, anything that causes it to fail will likely imply that the system is not functional. The router uses a dynamic label range from values 18,432 through 262,143 (131,071 in chassis modes lower than D) to allocate all dynamic labels, including RSVP and BGP allocated labels and VC labels.

### **5.2.3.3 LDP Configuration**

The router 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, LDP targeted sessions can be set up to specific endpoints. Targeted-session parameters are configurable.

### **5.2.3.4 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.

### **5.2.3.5 Service Manager**

All interaction occurs between LDP and the service manager, since LDP is used primarily to exchange labels for Layer 2 services. In this context, the service manager informs LDP when an LDP session is to be set up or torn down, and when labels are to be exchanged or withdrawn. In turn, LDP informs service manager of relevant LDP events, such as connection setups and failures, timeouts, labels signaled/withdrawn.

## **5.2.4 Execution Flow**

LDP activity in the operating system is limited to service-related signaling. Therefore, the configurable parameters are restricted to system-wide parameters, such as hello and keepalive timeouts.

### **5.2.4.1 Initialization**

LDP makes sure that the various prerequisites, such as ensuring the system IP interface is operational, the label manager is operational, and there is memory available, are met. It then allocates itself a pool of memory and initializes its databases.

### **5.2.4.2 Session Lifetime**

In order for a targeted LDP (T-LDP) session to be established, an adjacency must be created. The LDP extended discovery mechanism requires hello messages to be exchanged between two peers for session establishment. After the adjacency establishment, session setup is attempted.

#### **5.2.4.2.1 Adjacency Establishment**

In the router, the adjacency management is done through the establishment of a Service Distribution Path (SDP) object, which is a service entity in the Nokia service model.

The Nokia service model uses logical entities that interact to provide a service. The service model requires the service provider to create configurations for four main entities:

- Customers
- Services
- Service Access Paths (SAPs) on the local routers
- Service Distribution Points (SDPs) that connect to one or more remote routers.

An SDP is the network-side termination point for a tunnel to a remote router. An SDP defines a local entity that includes the system IP address of the remote routers and a path type. Each SDP comprises:

- The SDP ID
- The transport encapsulation type, either MPLS or GRE
- 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, since only one adjacency and one LDP session exists between the pair of nodes.

An SDP is a uni-directional object, so a pair of SDPs pointing at each other must be configured in order for an LDP adjacency to be established. Once an adjacency is established, it is maintained through periodic hello messages.

#### **5.2.4.2.2 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 liveliness.

Since TCP is back-pressured by the receiver, it is necessary to be able to push that back-pressure all the way into the protocol. Packets that cannot be sent are buffered on the session object and re-attempted as the back-pressure eases.

## 5.2.5 Label Exchange

Label exchange is initiated by the service manager. When an SDP is attached to a service (for example, the service gets a transport tunnel), a message is sent from the service manager to LDP. This causes a label mapping message to be sent. Additionally, when the SDP binding is removed from the service, the VC label is withdrawn. The peer must send a label release to confirm that the label is not in use.

### 5.2.5.1 Other Reasons for Label Actions

Other reasons for label actions include:

- MTU changes: LDP withdraws the previously assigned label, and re-signals the FEC with the new MTU in the interface parameter.
- Clear labels: When a service manager command is issued to clear the labels, the labels are withdrawn, and new label mappings are issued.
- SDP down: When an SDP goes administratively down, the VC label associated with that SDP for each service is withdrawn.
- Memory allocation failure: If there is no memory to store a received label, it is released.
- VC type unsupported: When an unsupported VC type is received, the received label is released.

### 5.2.5.2 Cleanup

LDP closes all sockets, frees all memory, and shuts down all its tasks when it is deleted, so its memory usage is 0 when it is not running.

### 5.2.5.3 Configuring Implicit Null Label

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

Enable the use of the implicit null option, for all LDP FECs for which this node is the egress LER, using the following command:

### **config>router>ldp>implicit-null-label**

When the user changes the implicit null configuration option, LDP withdraws all the FECs and re-advertises them using the new label value.

## **5.2.6 Global LDP Filters**

Both inbound and outbound LDP label binding filtering are supported.

Inbound filtering is performed by way of the configuration of an import policy to control the label bindings an LSR accepts from its peers. Label bindings can be filtered based on:

- Prefix-list: Match on bindings with the specified prefix/prefixes.
- Neighbor: Match on bindings received from the specified peer.

The default import policy is to accept all FECs received from peers.

Outbound filtering is performed by way of the configuration of an export policy. The Global LDP export policy can be used to explicitly originate label bindings for local interfaces. The Global LDP export policy does not filter out or stop propagation of any FEC received from neighbors. Use the LDP peer export prefix policy for this purpose. The system IP address AND static FECs cannot be blocked using an export policy.

Export policy enables configuration of a policy to advertise label bindings based on:

- Direct: All local subnets.
- Prefix-list: Match on bindings with the specified prefix or prefixes.

The default export policy is to originate label bindings for system address only and to propagate all FECs received from other LDP peers.

Finally, the 'neighbor interface' statement inside a global import policy is not considered by LDP.

### **5.2.6.1 Per LDP Peer FEC Import and Export Policies**

The FEC prefix export policy provides a way to control which FEC prefixes received from prefixes received from other LDP and T-LDP peers are re-distributed to this LDP peer.

The user configures the FEC prefix export policy using the following command:

---

```
config>router>ldp>session-params>peer>export-prefixes policy-name
```

By default, all FEC prefixes are exported to this peer.

The FEC prefix import policy provides a mean of controlling which FEC prefixes received from this LDP peer are imported and installed by LDP on this node. If resolved these FEC prefixes are then re-distributed to other LDP and T-LDP peers.

The user configures the FEC prefix export policy using the following command:

```
config>router>ldp>session-params>peer>import-prefixes policy-name
```

By default, all FEC prefixes are imported from this peer.

## 5.2.7 Configuring Multiple LDP LSR ID

The multiple LDP LSR-ID feature provides the ability to configure and initiate multiple Targeted LDP (T-LDP) sessions on the same system using different LDP LSR-IDs. In the current implementation, all T-LDP sessions must have the LSR-ID match the system interface address. This feature continues to allow the use of the system interface by default, but also any other network interface, including a loopback, address on a per T-LDP session basis. The LDP control plane will not allow more than a single T-LDP session with different local LSR ID values to the same LSR-ID in a remote node.

An SDP of type LDP can use a provisioned targeted session with the local LSR-ID set to any network IP for the T-LDP session to the peer matching the SDP far-end address. If, however, no targeted session has been explicitly pre-provisioned to the far-end node under LDP, then the SDP will auto-establish one but will use the system interface address as the local LSR-ID.

An SDP of type RSVP must use an RSVP LSP with the destination address matching the remote node LDP LSR-ID. An SDP of type GRE can only use a T-LDP session with a local LSR-ID set to the system interface.

The multiple LDP LSR-ID feature also provides the ability to use the address of the local LDP interface, or any other network IP interface configured on the system, as the LSR-ID to establish link LDP Hello adjacency and LDP session with directly connected LDP peers. The network interface can be a loopback or not.

Link LDP sessions to all peers discovered over a given LDP interface share the same local LSR-ID. However, LDP sessions on different LDP interfaces can use different network interface addresses as their local LSR-ID.

By default, the link and targeted LDP sessions to a peer use the system interface address as the LSR-ID unless explicitly configured using this feature. The system interface must always be configured on the router or else the LDP protocol will not come up on the node. There is no requirement to include it in any routing protocol.

When an interface other than system is used as the LSR-ID, the transport connection (TCP) for the link or targeted LDP session will also use the address of that interface as the transport address.

## 5.2.8 T-LDP hello reduction

This feature implements a new mechanism to suppress the transmission of the Hello messages following the establishment of a Targeted LDP session between two LDP peers. The Hello adjacency of the targeted session does not require periodic transmission of Hello messages as in the case of a link LDP session. In link LDP, one or more peers can be discovered over a given network IP interface and as such, the periodic transmission of Hello messages is required to discover new peers in addition to the periodic Keep-Alive message transmission to maintain the existing LDP sessions. A Targeted LDP session is established to a single peer. Thus, once the Hello Adjacency is established and the LDP session is brought up over a TCP connection, Keep-Alive messages are sufficient to maintain the LDP session.

When this feature is enabled, the targeted Hello adjacency is brought up by advertising the Hold-Time value the user configured in the Hello timeout parameter for the targeted session. The LSR node will then start advertising an exponentially increasing Hold-Time value in the Hello message as soon as the targeted LDP session to the peer is up. Each new incremented Hold-Time value is sent in a number of Hello messages equal to the value of the Hello reduction factor before the next exponential value is advertised. This provides time for the two peers to settle on the new value. When the Hold-Time reaches the maximum value of 0xffff (binary 65535), the two peers will send Hello messages at a frequency of every  $[(65535-1)/\text{local helloFactor}]$  seconds for the lifetime of the targeted-LDP session (for example, if the local Hello Factor is three (3), then Hello messages will be sent every 21844 seconds).

Both LDP peers must be configured with this feature to bring gradually their advertised Hold-Time up to the maximum value. If one of the LDP peers does not, the frequency of the Hello messages of the targeted Hello adjacency will continue to be governed by the smaller of the two Hold-Time values. This feature complies to *draft-pdutta-mpls-tldp-hello-reduce*.

## 5.2.9 Tracking a T-LDP Peer with BFD

BFD tracking of an LDP session associated with a T-LDP adjacency allows for faster detection of the liveliness of the session by registering the peer transport address of a LDP session with a BFD session. The source or destination address of the BFD session is the local or remote transport address of the targeted or link (if peers are directly connected) Hello adjacency which triggered the LDP session.

By enabling BFD for a selected targeted session, the state of that session is tied to the state of the underneath BFD session between the two nodes. The parameters used for the BFD are set with the BFD command under the IP interface which has the source address of the TCP connection.

## 5.2.10 Link LDP Hello Adjacency Tracking with BFD

LDP can only track an LDP peer using the Hello and Keep-Alive timers. If an IGP protocol registered with BFD on an IP interface to track a neighbor, and the BFD session times out, the next-hop for prefixes advertised by the neighbor are no longer resolved. This however does not bring down the link LDP session to the peer since the LDP peer is not directly tracked by BFD.

In order to properly track the link LDP peer, LDP needs to track the Hello adjacency to its peer by registering with BFD.

The user effects Hello adjacency tracking with BFD by enabling BFD on an LDP interface:

```
config>router>ldp>if-params>if>enable-bfd [ipv4][ipv6]
```

The parameters used for the BFD session, i.e., transmit-interval, receive-interval, and multiplier, are those configured under the IP interface:

```
config>router>if>bfd
```

The source or destination address of the BFD session is the local or remote address of link Hello adjacency. When multiple links exist to the same LDP peer, a Hello adjacency is established over each link. However, a single LDP session will exist to the peer and will use a TCP connection over one of the link interfaces. Also, a separate BFD session should be enabled on each LDP interface. If a BFD session times out on a specific link, LDP will immediately bring down the Hello adjacency on that link. In addition, if there are FECs that have their primary NHLFE over this



link, LDP triggers the LDP FRR procedures by sending to IOM and line cards the neighbor/next-hop down message. This will result in moving the traffic of the impacted FECs to an LFA next-hop on a different link to the same LDP peer or to an LFA backup next-hop on a different LDP peer depending on the lowest backup cost path selected by the IGP SPF.

As soon as the last Hello adjacency goes down as a result of the BFD timing out, the LDP session goes down and the LDP FRR procedures will be triggered. This will result in moving the traffic to an LFA backup next-hop on a different LDP peer.

### 5.2.11 LDP LSP Statistics

RSVP-TE LSP statistics is extended to LDP to provide the following counters:

- 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

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

### 5.2.12 MPLS Entropy Label

The router supports the MPLS entropy label (RFC 6790) on LDP LSPs used for IGP and BGP shortcuts. This allows LSR nodes in a network to load-balance labeled packets in a much more granular fashion than allowed by simply hashing on the standard label stack.

## 5.3 TTL Security for BGP and LDP

The BGP TTL Security Hack (BTSH) was originally designed to protect the BGP infrastructure from CPU utilization-based attacks. It is derived from the fact that the vast majority of ISP eBGP peerings are established between adjacent routers. Since TTL spoofing is considered nearly impossible, a mechanism based on an expected TTL value can provide a simple and reasonably robust defense from infrastructure attacks based on forged BGP packets.

While TTL Security Hack (TSH) is most effective in protecting directly connected peers, it can also provide a lower level of protection to multi-hop sessions. When a multi-hop BGP session is required, the expected TTL value can be set to 255 minus the configured range-of-hops. This approach can provide a qualitatively lower degree of security for BGP (such as a DoS attack could, theoretically, be launched by compromising a box in the path). However, BTSH will catch a vast majority of observed distributed DoS (DDoS) attacks against eBGP.

TSH can be used to protect LDP peering sessions as well. For details, see draft-chen-ldp-ttl-xx.txt, *TTL-Based Security Option for LDP Hello Message*.

The TSH implementation supports the ability to configure TTL security per BGP/LDP peer and evaluate (in hardware) the incoming TTL value against the configured TTL value. If the incoming TTL value is less than the configured TTL value, the packets are discarded and a log is generated.

## 5.4 ECMP Support for LDP

ECMP support for LDP performs load balancing for LDP based LSPs by having multiple outgoing next-hops for a given IP prefix on ingress and transit LSRs.

An LSR that has multiple equal cost paths to a given IP prefix can receive an LDP label mapping for this prefix from each of the downstream next-hop peers. As the LDP implementation uses the liberal label retention mode, it retains all the labels for an IP prefix received from multiple next-hop peers.

Without ECMP support for LDP, only one of these next-hop peers will be selected and installed in the forwarding plane. The algorithm used to determine the next-hop peer to be selected involves looking up the route information obtained from the RTM for this prefix and finding the first valid LDP next-hop peer (for example, the first neighbor in the RTM entry from which a label mapping was received). If, for some reason, the outgoing label to the installed next-hop is no longer valid, say the session to the peer is lost or the peer withdraws the label, a new valid LDP next-hop peer will be selected out of the existing next-hop peers and LDP will reprogram the forwarding plane to use the label sent by this peer.

With ECMP support, all the valid LDP next-hop peers, those that sent a label mapping for a given IP prefix, will be installed in the forwarding plane. In both cases, ingress LER and transit LSR, an ingress label will be mapped to the nexthops that are in the RTM and from which a valid mapping label has been received. The forwarding plane will then use an internal hashing algorithm to determine how the traffic will be distributed amongst these multiple next-hops, assigning each “flow” to a particular next-hop.

The hash algorithm at LER and transit LSR are described in the LAG and ECMP Hashing section of the SR OS Interface Guide.

## 5.4.1 Label Operations

If an LSR is the ingress for a given IP prefix, LDP programs a push operation for the prefix in the forwarding engine. This creates an LSP ID to the Next Hop Label Forwarding Entry (NHLFE) (LTN) mapping and an LDP tunnel entry in the forwarding plane. LDP will also inform the Tunnel Table Manager (TTM) of this tunnel. Both the LTN entry and the tunnel entry will have a NHLFE for the label mapping that the LSR received from each of its next-hop peers.

If the LSR is to behave as a transit for a given IP prefix, LDP will program a swap operation for the prefix in the forwarding engine. This involves creating an Incoming Label Map (ILM) entry in the forwarding plane. The ILM entry will have to map an incoming label to possibly multiple NHLFEs. If an LSR is an egress for a given IP prefix, LDP will program a POP entry in the forwarding engine. This too will result 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 will consult the LTN entry and will use a hashing algorithm to map the packet to one of the NHLFEs (push label) and forward the packet to the corresponding next-hop peer. For labeled packets arriving at a transit or egress LSR, the forwarding plane will consult the ILM entry and either use a hashing algorithm to map it to one of the NHLFEs if they exist (swap label) or simply route the packet if there are no NHLFEs (pop label).

Static FEC swap will not be activated unless there is a matching route in system route table that also matches the user configured static FEC next-hop.

## 5.5 Unnumbered Interface Support in LDP

This feature allows LDP to establish Hello adjacency and to resolve unicast and multicast FECs over unnumbered LDP interfaces.

This feature also extends the support of lsp-ping, p2mp-lsp-ping, and ldp-treetrace to test an LDP unicast or multicast FEC which is resolved over an unnumbered LDP interface.

## 5.5.1 Feature Configuration


This feature does not introduce a new CLI command for adding an unnumbered interface into LDP. Rather, the **fec-originate** command is extended to specify the interface name because an unnumbered interface does not have an IP address of its own. The user can, however, specify the interface name for numbered interfaces.

See the CLI section for the changes to the **fec-originate** command.

## 5.5.2 Operation of LDP over an Unnumbered IP Interface

Consider the setup shown in [Figure 57](#).

**Figure 57** LDP Adjacency and Session over Unnumbered Interface



LSR A and LSR B have the following LDP identifiers respectively:

<LSR Id=A> : <label space id=0>

<LSR Id=B> : <label space id=0>

There are two P2P unnumbered interfaces between LSR A and LSR B. These interfaces are identified on each system with their unique local link identifier. In other words, the combination of {Router-ID, Local Link Identifier} uniquely identifies the interface in OSPF or IS-IS throughout the network.

A borrowed IP address is also assigned to the interface to be used as the source address of IP packets which need to be originated from the interface. The borrowed IP address defaults to the system loopback interface address, A and B respectively in this setup. The user can change the borrowed IP interface to any configured IP interface, loopback or not, by applying the following command:

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

When the unnumbered interface is added into LDP, it will have the following behavior.

### 5.5.2.1 Link LDP

Hello adjacency will be brought up using link Hello packet with source IP address set to the interface borrowed IP address and a destination IP address set to 224.0.0.2.

As a consequence of (1), Hello packets with the same source IP address should be accepted when received over parallel unnumbered interfaces from the same peer LSR-ID. The corresponding Hello adjacencies would be associated with a single LDP session.

The transport address for the TCP connection, which is encoded in the Hello packet, will always be set to the LSR-ID of the node regardless if the user enabled the interface option under **config>router>ldp>if-params>if>ipv4>transport-address**.

The user can configure the local-lsr-id option on the interface and change the value of the LSR-ID to either the local interface or to some other interface name, loopback or not, numbered or not. If the local interface is selected or the provided interface name corresponds to an unnumbered IP interface, the unnumbered interface borrowed IP address will be used as the LSR-ID. In all cases, the transport address for the LDP session will be updated to the new LSR-ID value but the link Hello packets will continue to use the interface borrowed IP address as the source IP address.

The LSR with the highest transport address, i.e., LSR-ID in this case, will bootstrap the TCP connection and LDP session.

Source and destination IP addresses of LDP packets are the transport addresses, i.e., LDP LSR-IDs of systems A and B in this case.

### 5.5.2.2 Targeted LDP

Source and destination addresses of targeted Hello packet are the LDP LSR-IDs of systems A and B.

The user can configure the local-lsr-id option on the targeted session and change the value of the LSR-ID to either the local interface or to some other interface name, loopback or not, numbered or not. If the local interface is selected or the provided interface name corresponds to an unnumbered IP interface, the unnumbered interface borrowed IP address will be used as the LSR-ID. In all cases, the transport address for the LDP session and the source IP address of targeted Hello message will be updated to the new LSR-ID value.

The LSR with the highest transport address, i.e., LSR-ID in this case, will bootstrap the TCP connection and LDP session.

Source and destination IP addresses of LDP messages are the transport addresses, i.e., LDP LSR-IDs of systems A and B in this case.

### 5.5.2.3 FEC Resolution

LDP will advertise/withdraw unnumbered interfaces using the Address/Address-Withdraw message. The borrowed IP address of the interface is used.

A FEC can be resolved to an unnumbered interface in the same way as it is resolved to a numbered interface. The outgoing interface and next-hop are looked up in RTM cache. The next-hop consists of the router-id and link identifier of the interface at the peer LSR.

LDP FEC ECMP next-hops over a mix of unnumbered and numbered interfaces is supported.

All LDP FEC types are supported.

The **fec-originate** command is supported when the next-hop is over an unnumbered interface.

All LDP features are supported except for the following:


- BFD cannot be enabled on an unnumbered LDP interface. This is a consequence of the fact that BFD is not supported on unnumbered IP interface on the system.
- As a consequence of (1), LDP FRR procedures will not be triggered via a BFD session timeout but only by physical failures and local interface down events.
- Unnumbered IP interfaces cannot be added into LDP global and peer prefix policies.

## 5.6 LDP over RSVP Tunnels

LDP over RSVP-TE provides end-to-end tunnels that have two important properties, fast reroute and traffic engineering which are not available in LDP. LDP over RSVP-TE is focused at large networks (over 100 nodes in the network). Simply using end-to-end RSVP-TE tunnels will not scale. While an LER may not have that many tunnels, any transit node will potentially have thousands of LSPs, and if each transit node also has to deal with detours or bypass tunnels, this number can make the LSR overly burdened.

LDP over RSVP-TE allows tunneling of user packets using an LDP LSP inside an RSVP LSP. The main application of this feature is for deployment of MPLS based services, for example, VPRN, VLL, and VPLS services, in large scale networks across multiple IGP areas without requiring full mesh of RSVP LSPs between PE routers.

**Figure 58 LDP over RSVP Application**



The network displayed in [Figure 58](#) consists of two metro areas, Area 1 and 2 respectively, and a core area, Area 3. Each area makes use of TE LSPs to provide connectivity between the edge routers. In order to enable services between PE1 and PE2 across the three areas, LSP1, LSP2, and LSP3 are set up using RSVP-TE. There are in fact 6 LSPs required for bidirectional operation but we will refer to each bi-directional LSP with a single name, for example, LSP1. A targeted LDP (T-LDP) session is associated with each of these bidirectional LSP tunnels. That is, a T-LDP adjacency is created between PE1 and ABR1 and is associated with LSP1 at each end. The same is done for the LSP tunnel between ABR1 and ABR2, and finally between ABR2 and PE2. The loopback address of each of these routers is advertised using T-LDP. Similarly, backup bidirectional LDP over RSVP tunnels, LSP1a and LSP2a, are configured by way of ABR3.

This setup effectively creates an end-to-end LDP connectivity which can be used by all PEs to provision services. The RSVP LSPs are used as a transport vehicle to carry the LDP packets from one area to another. Only the user packets are tunneled over the RSVP LSPs. The T-LDP control messages are still sent unlabeled using the IGP shortest path.

In this application, the bi-directional RSVP LSP tunnels are not treated as IP interfaces and are not advertised back into the IGP. A PE must always rely on the IGP to look up the next hop for a service packet. LDP-over-RSVP introduces a new tunnel type, tunnel-in-tunnel, in addition to the existing LDP tunnel and RSVP tunnel types. If multiple tunnels types match the destination PE FEC lookup, LDP will prefer an LDP tunnel over an LDP-over-RSVP tunnel by default.


The design in [Figure 58](#) allows a service provider to build and expand each area independently without requiring a full mesh of RSVP LSPs between PEs across the three areas.

To participate in a VPRN service, the PE1 and PE2 perform the autobind to LDP. The LDP label which represents the target PE loopback address is used below the RSVP LSP label. Therefore a 3 label stack is required.

In order to provide a VLL service, PE1 and PE2 are still required to set up a targeted LDP session directly between them. Again a 3 label stack is required, the RSVP LSP label, followed by the LDP label for the loopback address of the destination PE, and finally the pseudowire label (VC label).

This implementation supports a variation of the application in [Figure 58](#), in which area 1 is an LDP area. In that case, PE1 will push a two label stack while ABR1 will swap the LDP label and push the RSVP label as illustrated in [Figure 59](#). LDP-over-RSVP tunnels can also be used as IGP shortcuts.

**Figure 59 LDP over RSVP Application Variant**



## 5.6.1 Signaling and Operation

- [LDP Label Distribution and FEC Resolution](#)
- [Default FEC Resolution Procedure](#)



### 5.6.1.1 LDP Label Distribution and FEC Resolution

The user creates a targeted LDP (T-LDP) session to an ABR or the destination PE. This results in LDP hellos being sent between the two routers. These messages are sent unlabeled over the IGP path. Next, the user enables LDP tunneling on this T-LDP session and optionally specifies a list of LSP names to associate with this T-LDP session. By default, all RSVP LSPs which terminate on the T-LDP peer are candidates for LDP-over-RSVP tunnels. At this point in time, the LDP FECs resolving to RSVP LSPs are added into the Tunnel Table Manager as tunnel-in-tunnel type.

If LDP is running on regular interfaces also, the prefixes LDP learns are going to be distributed over both the T-LDP session as well as regular IGP interfaces. However, only /32 FEC prefixes will be resolved over RSVP LSPs. The policy controls which prefixes go over the T-LDP session, for example, only /32 prefixes, or a particular prefix range.

LDP-over-RSVP works with both OSPF and ISIS. These protocols include the advertising router when adding an entry to the RTM. LDP-over-RSVP tunnels can be used as shortcuts for BGP next-hop resolution.

### 5.6.1.2 Default FEC Resolution Procedure

When LDP tries to resolve a prefix received over a T-LDP session, it performs a lookup in the Routing Table Manager (RTM). This lookup returns the next hop to the destination PE and the advertising router (ABR or destination PE itself). If the next-hop router advertised the same FEC over link-level LDP, LDP will prefer the LDP tunnel by default unless the user explicitly changed the default preference using the system wide prefer-tunnel-in-tunnel command. If the LDP tunnel becomes unavailable, LDP will select an LDP-over-RSVP tunnel if available.

When searching for an LDP-over-RSVP tunnel, LDP selects the advertising router(s) with best route. If the advertising router matches the T-LDP peer, LDP then performs a second lookup for the advertising router in the Tunnel Table Manager (TTM) which returns the user configured RSVP LSP with the best metric. If there are more than one configured LSP with the best metric, LDP selects the first available LSP.

If all user configured RSVP LSPs are down, no more action is taken. If the user did not configure any LSPs under the T-LDP session, the lookup in TTM will return the first available RSVP LSP which terminates on the advertising router with the lowest metric.

### 5.6.1.3 FEC Resolution Procedure When prefer-tunnel-in-tunnel is Enabled

When LDP tries to resolve a prefix received over a T-LDP session, it performs a lookup in the Routing Table Manager (RTM). This lookup returns the next hop to the destination PE and the advertising router (ABR or destination PE itself).

When searching for an LDP-over-RSVP tunnel, LDP selects the advertising router(s) with best route. If the advertising router matches the targeted LDP peer, LDP then performs a second lookup for the advertising router in the Tunnel Table Manager (TTM) which returns the user configured RSVP LSP with the best metric. If there are more than one configured LSP with the best metric, LDP selects the first available LSP.

If all user configured RSVP LSPs are down, then an LDP tunnel will be selected if available.

If the user did not configure any LSPs under the T-LDP session, a lookup in TTM will return the first available RSVP LSP which terminates on the advertising router. If none are available, then an LDP tunnel will be selected if available.

## 5.6.2 Rerouting Around Failures

Every failure in the network can be protected against, except for the ingress and egress PEs. All other constructs have protection available. These constructs are LDP-over-RSVP tunnel and ABR.

- [LDP-over-RSVP Tunnel Protection](#)
- [ABR Protection](#)

### 5.6.2.1 LDP-over-RSVP Tunnel Protection

An RSVP LSP can deal with a failure in two ways:

- If the LSP is a loosely routed LSP, then RSVP will find a new IGP path around the failure, and traffic will follow this new path. This may involve some churn in the network if the LSP comes down and then gets re-routed. The tunnel damping feature was implemented on the LSP so that all the dependent protocols and applications do not flap unnecessarily.

- If the LSP is a CSPF-computed LSP with the fast reroute option enabled, then RSVP will switch to the detour path very quickly. From that point, a new LSP will be attempted from the head-end (global revertive). When the new LSP is in place, the traffic switches over to the new LSP with make-before-break.


### 5.6.2.2 ABR Protection

If an ABR fails, then routing around the ABR requires that a new next-hop LDP-over-RSVP tunnel be found to a backup ABR. If an ABR fails, then the T-LDP adjacency fails. Eventually, the backup ABR becomes the new next hop (after SPF converges), and LDP learns of the new next-hop and can reprogram the new path.

## 5.7 LDP over RSVP Without Area Boundary

The LDP over RSVP capability set includes the ability to stitch LDP-over-RSVP tunnels at internal (non-ABR) OSPF and IS-IS routers.

**Figure 60 LDP over RSVP Without ABR Stitching Point**



al\_0214

In [Figure 60](#), assume that the user wants to use LDP over RSVP between router A and destination “Dest”. The first thing that happens is that either OSPF or IS-IS will perform an SPF calculation resulting in an SPF tree. This tree specifies the lowest possible cost to the destination. In the example shown, the destination “Dest” is reachable at the lowest cost through router X. The SPF tree will have the following path: A>C>E>G>X.

Using this SPF tree, router A will search for the endpoint that is closest (farthest/highest cost from the origin) to “Dest” that is eligible. Assuming that all LSPs in the above diagram are eligible, LSP endpoint G will be selected as it terminates on router G while other LSPs only reach routers C and E, respectively.

IGP and LSP metrics associated with the various LSP are ignored; only tunnel endpoint matters to IGP. The endpoint that terminates closest to “Dest” (highest IGP path cost) will be selected for further selection of the LDP over RSVP tunnels to that endpoint. The explicit path the tunnel takes may not match the IGP path that the SPF computes.

If router A and G have an additional LSP terminating on router G, there would now be two tunnels both terminating on the same router closest to the final destination. For IGP, it does not make any difference on the numbers of LSPs to G, only that there is at least one LSP to G. In this case, the LSP metric will be considered by LDP when deciding which LSP to stitch for the LDP over RSVP connection.

The IGP only passes endpoint information to LDP. LDP looks up the tunnel table for all tunnels to that endpoint and picks up the one with the least tunnel metric. There may be many tunnels with the same least cost. Only /32 FEC prefixes will be resolved over RSVP LSPs within an area.

### 5.7.1 LDP over RSVP and ECMP

ECMP for LDP over RSVP is supported (also see [ECMP Support for LDP](#)). If ECMP applies, all LSP endpoints found over the ECMP IGP path will be installed in the routing table by the IGP for consideration by LDP. IGP costs to each endpoint may differ because IGP selects the farthest endpoint per ECMP path.

LDP will choose the endpoint that is highest cost in the route entry and will do further tunnel selection over those endpoints. If there are multiple endpoints with equal highest cost, then LDP will consider all of them.

---

## 5.8 Class-Based Forwarding of LDP Prefix Packets over IGP Shortcuts

Within large ISP networks, services are typically required from any PE to any PE and can traverse multiple domains. Also, within a service, different traffic classes can co-exist, each with specific requirements on latency and jitter.

The class-based forwarding feature enables service providers to control which LSPs, of a set of ECMP tunnel next-hops that resolve an LDP FEC prefix, to forward packets that were classified to specific forwarding classes, as opposed to normal ECMP spraying where packets are sprayed over the whole set of LSPs.

### 5.8.1 Configuration and Operation

To achieve the behavior described above, the user must first enable the following:

- IGP shortcuts or forwarding adjacencies in the routing instance
- ECMP
- the advertisement of unicast prefix FECs on the Targeted LDP session to the peer
- class-based forwarding in the LDP context

Enabling these options is achieved by using the following commands:

Either one of:

- **config>router>isis>rsvp-shortcut**
- **config>router>ospf>rsvp-shortcut**
- Or one of:
  - **config>router>isis>advertise-tunnel-link**
  - **config>router>ospf>advertise-tunnel-link**
- All of:
  - **config>router>ecmp *max-ecmp-routes***
  - **config>router>ldp>targ-session>peer>tunneling**
  - **config>router>ldp>class-forwarding**

If the user specifies LSP names under the **tunneling** option, these LSPs are not directly used by LDP when the **rsvp-shortcut** option is enabled. With IGP shortcuts, the set of tunnel next-hops is always provided by IGP in RTM. Consequently, the class-based forwarding rules described below do not apply to this set of named LSPs unless they were populated by IGP in RTM as next-hops for a prefix.

The **prefer-tunnel-in-tunnel** must be disabled for class-based forwarding to apply to LDP prefixes which are the endpoint of the tunnels.

The user must also bind traffic classes to designated LSPs. This is performed using the following commands:

```
config>router>mpls>lsp>class-forwarding>fc {be | l2 | af | l1 | h2 | ef | h1 | nc}
```

The user can also designate a given LSP as a Default LSP using the following command:

```
config>router>mpls>lsp>class-forwarding>default-lsp
```

These two commands can also be passed in the **lsp-template** context such that LSPs created from that template will have the assigned Class-Based Forwarding (CBF) configurations.

When an LDP prefix is resolved to a set of ECMP tunnel next hops, the selection process by which the set is returned does not take into account any CBF configuration. As such, even if the user has assigned CBF configurations to one or more LSPs, those may not be selected as part of the set of ECMP tunnel next hops. The assignments of CBF configurations are done on a per-LSP (or LSP template) basis and, as such, are independent one from another. The evaluation of the consistency of the assignments is performed by LDP at the time the FEC is resolved to a set of ECMP tunnel next hops, and the following rules are applied.

- If no single LSP of the set has a CBF configuration assigned (either a forwarding class or the **default-lsp** option), then normal ECMP spraying will occur over the whole set of LSPs.
- If at least one LSP has a CBF configuration assigned, then class-based forwarding will occur. If the default-lsp option has not been assigned to an LSP, one will be automatically selected for that assignment by LDP. That LSP is the one with the lowest tunnel-id amongst the set of LSPs with one (or more) forwarding classes assigned to.
- Multiple LSPs can have the same forwarding class assigned. However, for each of these forwarding classes, only a single LSP will be used to forward packets classified into this forwarding class. That LSP is the one with the lowest tunnel-id amongst those sharing a given forwarding class.

- Similarly, multiple LSPs can have the **default-lsp** configuration assigned. Only a single one will be designated to be the Default LSP. That LSP is the one with the lowest tunnel-id amongst those with the **default-lsp** option assigned.

Therefore, under normal conditions, LDP prefix packets will be sprayed over a set of ECMP tunnel next-hops by selecting either the LSP to which is assigned the forwarding class of the packets, if one exists, or the Default LSP, if one does not exist. However, the CBF is suspended until LDP downloads a new consistent set of tunnel next-hops for the FEC. For example, if the IOM detects that the LSP to which is assigned a forwarding class is not usable, it will switch the forwarding of packets classified to that forwarding class into the Default LSP, and if the IOM detects that the Default LSP is not usable, then it will revert to regular ECMP spraying across all tunnels in the set of ECMP tunnel next-hops.

In case a user changes (adds, modifies, or deletes) the CBF configuration associated to an LSP which has previously been selected as part of a set of ECMP tunnel next hops, this change will automatically lead to an updated FEC resolution and CBF consistency check and may lead to an update of the forwarding configuration.

This functionality only applies to LSR forwarding LDP FEC prefix packets over a set of MPLS LSPs using IGP shortcuts. It does not apply to LER forwarding of shortcut packets over LDP FEC which is resolved to a set of MPLS LSPs using IGP shortcuts, nor does it apply to LER forwarding of packets of VPRN and Layer-2 services, which use auto-binding to LDP when the LDP FEC is resolved to a set of MPLS LSPs using IGP shortcuts.

## 5.9 LDP ECMP Uniform Failover

LDP ECMP uniform failover allows the fast re-distribution by the ingress data path of packets forwarded over an LDP FEC next-hop to other next-hops of the same FEC when the currently used next-hop fails. The switchover is performed within a bounded time, which does not depend on the number of impacted LDP ILMs (LSR role) or service records (ingress LER role). The uniform failover time is only supported for a single LDP interface or LDP next-hop failure event.

This feature complements the coverage provided by the LDP Fast-ReRoute (FRR) feature, which provides a Loop-Free Alternate (LFA) backup next-hop with uniform failover time. Prefixes that have one or more ECMP next-hop protection are not programmed with a LFA back-up next-hop, and vice-versa.

The LDP ECMP uniform failover feature builds on the concept of Protect Group ID (PG-ID) introduced in LDP FRR. LDP assigns a unique PG-ID to all FECs that have their primary Next-Hop Label Forwarding Entry (NHLFE) resolved to the same outgoing interface and next-hop.

When an ILM record (LSR role) or LSPid-to-NHLFE (LTN) record (LER role) is created on the IOM, it has the PG-ID of each ECMP NHLFE the FEC is using.

When a packet is received on this ILM/LTN, the hash routine selects one of the up to 32, or the ECMP value configured on the system, whichever is less, ECMP NHLFEs for the FEC based on a hash of the packet's header. If the selected NHLFE has its PG-ID in DOWN state, the hash routine re-computes the hash to select a backup NHLFE among the first 16, or the ECMP value configured on the system, whichever is less, NHLFEs of the FEC, excluding the one that is in DOWN state. Packets of the subset of flows that resolved to the failed NHLFE are thus sprayed among a maximum of 16 NHLFEs.

LDP then re-computes the new ECMP set to exclude the failed path and downloads it into the IOM. At that point, the hash routine will update the computation and begin spraying over the updated set of NHLFEs.

LDP sends the DOWN state update of the PG-ID to the IOM when the outgoing interface or a specific LDP next-hop goes down. This can be the result of any of the following events:

- Interface failure detected directly.
- Failure of the LDP session detected via T-LDP BFD or LDP Keep-Alive.
- Failure of LDP Hello adjacency detected via link LDP BFD or LDP Hello.

In addition, PIP will send an interface down event to the IOM if the interface failure is detected by other means than the LDP control plane or BFD. In that case, all PG-IDs associated with this interface will have their state updated by the IOM.

When tunneling LDP packets over an RSVP LSP, it is the detection of the T-LDP session going down, via BFD or Keep-Alive, which triggers the LDP ECMP uniform failover procedures. If the RSVP LSP alone fails and the latter is not protected by RSVP FRR, the failure event will trigger the re-resolution of the impacted FECs in the slow path.

When a multicast LDP (mLDP) FEC is resolved over ECMP links to the same downstream LDP LSR, the PG-ID DOWN state will cause packets of the FEC resolved to the failed link to be switched to another link using the linear FRR switchover procedures.

The LDP ECMP uniform failover is not supported in the following forwarding contexts:



- VPLS BUM packets.
- Packets forwarded to an IES/VP RN spoke-interface.
- Packets forwarded towards VPLS spoke in routed VPLS.

Finally, the LDP ECMP uniform failover is only supported for a single LDP interface, LDP next-hop, or peer failure event.

## 5.10 LDP Fast-Reroute for IS-IS and OSPF Prefixes

LDP Fast Re-Route (FRR) is a feature which allows the user to provide local protection for an LDP FEC by pre-computing and downloading to the IOM or XCM both a primary and a backup NHLFE for this FEC.

The primary NHLFE corresponds to the label of the FEC received from the primary next-hop as per standard LDP resolution of the FEC prefix in RTM. The backup NHLFE corresponds to the label received for the same FEC from a Loop-Free Alternate (LFA) next-hop.

The LFA next-hop pre-computation by IGP is described in RFC 5286 – “Basic Specification for IP Fast Reroute: Loop-Free Alternates”. LDP FRR relies on using the label-FEC binding received from the LFA next-hop to forward traffic for a given prefix as soon as the primary next-hop is not available. This means that a node resumes forwarding LDP packets to a destination prefix without waiting for the routing convergence. The label-FEC binding is received from the loop-free alternate next-hop ahead of time and is stored in the Label Information Base since LDP on the router operates in the liberal retention mode.

This feature requires that IGP performs the Shortest Path First (SPF) computation of an LFA next-hop, in addition to the primary next-hop, for all prefixes used by LDP to resolve FECs. IGP also populates both routes in the Routing Table Manager (RTM).

### 5.10.1 LDP FRR Configuration

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

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

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

Next the user enables the use by LDP of the LFA next-hop by configuring the following option:

**config>router>ldp>fast-reroute**

When this command is enabled, LDP will use both the primary next-hop and LFA next-hop, when available, for resolving the next-hop of an LDP FEC against the corresponding prefix in the RTM. This will result in LDP programming a primary NHLFE and a backup NHLFE into the IOM or XCM for each next-hop of a FEC prefix for the purpose of forwarding packets over the LDP FEC.

Because LDP can detect the loss of a neighbor/next-hop independently, it is possible that it switches to the LFA next-hop while IGP is still using the primary next-hop. In order to avoid this situation, it is recommended to enable IGP-LDP synchronization on the LDP interface:

**config>router>interface>ldp-sync-timer *seconds***

### 5.10.1.1 Reducing the Scope of the LFA Calculation by SPF

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

**config>router>isis>level>loopfree-alternate-exclude**  
**config>router>ospf>area>loopfree-alternate-exclude**

If IGP shortcut are also enabled in LFA SPF, the LSPs with destination address in that IS-IS level or OSPF area are also not included in the LFA SPF calculation.

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

**config>router>isis>interface> loopfree-alternate-exclude**  
**config>router>ospf>area>interface> loopfree-alternate-exclude**

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

Finally, the user can apply the same above commands for an OSPF instance within a VPRN service:

```
config>service>vprn>ospf>area>loopfree-alternate-exclude
config>service>vprn>ospf>area>interface>loopfree-alternate-exclude
```

## 5.10.2 LDP FRR Procedures

The LDP FEC resolution when LDP FRR is not enabled operates as follows. When LDP receives a *FEC, label* binding for a prefix, it will resolve it by checking if the exact prefix, or a longest match prefix when the **aggregate-prefix-match option** is enabled in LDP, exists in the routing table and is resolved against a next-hop which is an address belonging to the LDP peer which advertised the binding, as identified by its LSR-id. When the next-hop is no longer available, LDP de-activates the FEC and de-programs the NHLFE in the data path. LDP will also immediately withdraw the labels it advertised for this FEC and deletes the ILM in the data path unless the user configured the **label-withdrawal-delay** option to delay this operation. Traffic that is received while the ILM is still in the data path is dropped. When routing computes and populates the routing table with a new next-hop for the prefix, LDP resolves again the FEC and programs the data path accordingly.

When LDP FRR is enabled and an LFA backup next-hop exists for the FEC prefix in RTM, or for the longest prefix the FEC prefix matches to when **aggregate-prefix-match** option is enabled in LDP, LDP will resolve the FEC as above but will program the data path with both a primary NHLFE and a backup NHLFE for each next-hop of the FEC.

In order perform a switchover to the backup NHLFE in the fast path, LDP follows the uniform FRR failover procedures which are also supported with RSVP FRR.

When any of the following events occurs, LDP instructs in the fast path the IOM on the line cards to enable the backup NHLFE for each FEC next-hop impacted by this event. The IOM line cards do that by simply flipping a single state bit associated with the failed interface or neighbor/next-hop:

1. An LDP interface goes operationally down, or is admin shutdown. In this case, LDP sends a neighbor/next-hop down message to the IOM line cards for each LDP peer it has adjacency with over this interface.
2. An LDP session to a peer went down as the result of the Hello or Keep-Alive timer expiring over a specific interface. In this case, LDP sends a neighbor/next-hop down message to the IOM line cards for this LDP peer only.
3. The TCP connection used by a link LDP session to a peer went down, due say to next-hop tracking of the LDP transport address in RTM, which brings down the LDP session. In this case, LDP sends a neighbor/next-hop down message to the IOM line cards for this LDP peer only.
4. A BFD session, enabled on a T-LDP session to a peer, times-out and as a result the link LDP session to the same peer and which uses the same TCP connection as the T-LDP session goes also down. In this case, LDP sends a neighbor/next-hop down message to the IOM line cards for this LDP peer only.
5. A BFD session enabled on the LDP interface to a directly connected peer, times-out and brings down the link LDP session to this peer. In this case, LDP sends a neighbor/next-hop down message to the IOM line cards for this LDP peer only. BFD support on LDP interfaces is a new feature introduced for faster tracking of link LDP peers.

The tunnel-down-dump-time option or the label-withdrawal-delay option, when enabled, does not cause the corresponding timer to be activated for a FEC as long as a backup NHLFE is still available.

### 5.10.2.1 ECMP Considerations

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

Also, when the system ECMP value is set to **ecmp=1** or to **no ecmp**, which translates to the same and is the default value, SPF can use the overflow ECMP links as LFA next-hops in these two cases.

### 5.10.2.2 LDP FRR and LDP Shortcut

When LDP FRR is enabled in LDP and the `ldp-shortcut` option is enabled in the router level, in transit IPv4 packets and specific CPM generated IPv4 control plane packets with a prefix resolving to the LDP shortcut are protected by the backup LDP NHLFE.

### 5.10.2.3 LDP FRR and LDP-over-RSVP

When LDP-over-RSVP is enabled, the RSVP LSP is modeled as an endpoint, i.e., the destination node of the LSP, and not as a link in the IGP SPF. Thus, it is not possible for IGP to compute a primary or alternate next-hop for a prefix which FEC path is tunneled over the RSVP LSP. Only LDP is aware of the FEC tunneling but it cannot determine on its own a loop-free backup path when it resolves the FEC to an RSVP LSP.

As a result, LDP does not activate the LFA next-hop it learned from RTM for a FEC prefix when the FEC is resolved to an RSVP LSP. LDP will activate the LFA next-hop as soon as the FEC is resolved to direct primary next-hop.

LDP FEC tunneled over an RSVP LSP due to enabling the LDP-over-RSVP feature will thus not support the LDP FRR procedures and will follow the slow path procedure of prior implementation.

When the user enables the **lfa-only** option for an RSVP LSP, as described in [Loop-Free Alternate Calculation in the Presence of IGP shortcuts](#), the LSP will not be used by LDP to tunnel an LDP FEC even when IGP shortcut is disabled but LDP-over-RSVP is enabled in IGP.

### 5.10.2.4 LDP FRR and RSVP Shortcut (IGP Shortcut)

When an RSVP LSP is used as a shortcut by IGP, it is included by SPF as a P2P link and can also be optionally advertised into the rest of the network by IGP. Thus the SPF is able of using a tunneled next-hop as the primary next-hop for a given prefix. LDP is also able of resolving a FEC to a tunneled next-hop when the IGP shortcut feature is enabled.

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

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


The LFA SPF is extended to use IGP shortcuts as LFA next-hops as explained in [Loop-Free Alternate Calculation in the Presence of IGP shortcuts](#).

### 5.10.3 IS-IS and OSPF Support for Loop-Free Alternate Calculation

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

Figure 61 illustrates a simple network topology with point-to-point (P2P) interfaces and highlights three routes to reach router R5 from router R1.

**Figure 61 Topology with Primary and LFA Routes**




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

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

Consider the case where the primary next-hop uses a broadcast interface as illustrated in Figure 62

**Figure 62 Example Topology with Broadcast Interfaces**



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

The following are the detailed rules for this criterion as provided in RFC 5286:

- **Rule 1:** Link-protect LFA backup next-hop (primary next-hop R1-R3 is a P2P interface):  
 $\text{Distance\_opt}(\text{R2}, \text{R5}) < \text{Distance\_opt}(\text{R2}, \text{R1}) + \text{Distance\_opt}(\text{R1}, \text{R5})$   
and,  
 $\text{Distance\_opt}(\text{R2}, \text{R5}) \geq \text{Distance\_opt}(\text{R2}, \text{R3}) + \text{Distance\_opt}(\text{R3}, \text{R5})$
- **Rule 2:** Node-protect LFA backup next-hop (primary next-hop R1-R3 is a P2P interface):  
 $\text{Distance\_opt}(\text{R2}, \text{R5}) < \text{Distance\_opt}(\text{R2}, \text{R1}) + \text{Distance\_opt}(\text{R1}, \text{R5})$   
and,  
 $\text{Distance\_opt}(\text{R2}, \text{R5}) < \text{Distance\_opt}(\text{R2}, \text{R3}) + \text{Distance\_opt}(\text{R3}, \text{R5})$
- **Rule 3:** Link-protect LFA backup next-hop (primary next-hop R1-R3 is a broadcast interface):  
 $\text{Distance\_opt}(\text{R2}, \text{R5}) < \text{Distance\_opt}(\text{R2}, \text{R1}) + \text{Distance\_opt}(\text{R1}, \text{R5})$   
and,  
 $\text{Distance\_opt}(\text{R2}, \text{R5}) < \text{Distance\_opt}(\text{R2}, \text{PN}) + \text{Distance\_opt}(\text{PN}, \text{R5})$   
where; PN stands for the R1-R3 link Pseudo-Node.

For the case of P2P interface, if SPF finds multiple LFA next-hops for a given primary next-hop, it follows the following selection algorithm:

1. It will pick the node-protect type in favor of the link-protect type.
2. If there is more than one LFA next-hop within the selected type, then it will pick one based on the least cost.
3. If more than one LFA next-hop with the same cost results from Step B, then SPF will select the first one. This is not a deterministic selection and will vary following each SPF calculation.

For the case of a broadcast interface, a node-protect LFA is not necessarily a link protect LFA if the path to the LFA next-hop goes over the same PN as the primary next-hop. Similarly, a link protect LFA may not guarantee link protection if it goes over the same PN as the primary next-hop.

The selection algorithm when SPF finds multiple LFA next-hops for a given primary next-hop is modified as follows:

1. The algorithm splits the LFA next-hops into two sets:
  - The first set consists of LFA next-hops which *do not* go over the PN used by primary next-hop.
  - The second set consists of LFA next-hops which *do* go over the PN used by the primary next-hop.
2. If there is more than one LFA next-hop in the first set, it will pick the node-protect type in favor of the link-protect type.
3. If there is more than one LFA next-hop within the selected type, then it will pick one based on the least cost.
4. If more than one LFA next-hop with equal cost results from Step C, SPF will select the first one from the remaining set. This is not a deterministic selection and will vary following each SPF calculation.
5. If no LFA next-hop results from Step D, SPF will rerun Steps B-D using the second set.

This algorithm is more flexible than strictly applying Rule 3 above; the link protect rule in the presence of a PN and specified in RFC 5286. A node-protect LFA which does not avoid the PN; does not guarantee link protection, can still be selected as a last resort. The same thing, a link-protect LFA which does not avoid the PN may still be selected as a last resort. Both the computed primary next-hop and LFA next-hop for a given prefix are programmed into RTM.



### 5.10.3.1 Loop-Free Alternate Calculation in the Presence of IGP shortcuts

In order to expand the coverage of the LFA backup protection in a network, RSVP LSP based IGP shortcuts can be placed selectively in parts of the network and be used as an LFA backup next-hop.

When IGP shortcut is enabled in IS-IS or OSPF on a given node, all RSVP LSP originating on this node and with a destination address matching the router-id of any other node in the network are included in the main SPF by default.

In order to limit the time it takes to compute the LFA SPF, the user must explicitly enable the use of an IGP shortcut as LFA backup next-hop using one of a couple of new optional argument for the existing LSP level IGP shortcut command:

**config>router>mpls>lsp>igp-shortcut [lfa-protect | lfa-only]**

The **lfa-protect** option allows an LSP to be included in both the main SPF and the LFA SPFs. 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 computation selected a tunneled primary next-hop for a prefix, the LFA SPF will not select an LFA next-hop for this prefix and the protection of this prefix will rely on the RSVP LSP FRR protection. If the main SPF computation selected a direct primary next-hop, then the LFA SPF will select an LFA next-hop for this prefix but will prefer a direct LFA next-hop over a tunneled LFA next-hop.

The **lfa-only** option allows an LSP to be included in the LFA SPFs only such that the introduction of IGP shortcuts does not impact the main SPF decision. For a given prefix, the main SPF always selects a direct primary next-hop. The LFA SPF will select a an LFA next-hop for this prefix but will prefer a direct LFA next-hop over a tunneled LFA next-hop.

Thus the selection algorithm when SPF finds multiple LFA next-hops for a given primary next-hop is modified as follows:

1. The algorithm splits the LFA next-hops into two sets:
  - the first set consists of direct LFA next-hops
  - the second set consists of tunneled LFA next-hops. after excluding the LSPs which use the same outgoing interface as the primary next-hop.
2. The algorithms continues with first set if not empty, otherwise it continues with second set.
3. If the second set is used, the algorithm selects the tunneled LFA next-hop which endpoint corresponds to the node advertising the prefix.

- If more than one tunneled next-hop exists, it selects the one with the lowest LSP metric.
  - If still more than one tunneled next-hop exists, it selects the one with the lowest tunnel-id.
  - If none is available, it continues with rest of the tunneled LFAs in second set.
4. Within the selected set, the algorithm splits the LFA next-hops into two sets:
    - The first set consists of LFA next-hops which do not go over the PN used by primary next-hop.
    - The second set consists of LFA next-hops which go over the PN used by the primary next-hop.
  5. If there is more than one LFA next-hop in the selected set, it will pick the node-protect type in favor of the link-protect type.
  6. If there is more than one LFA next-hop within the selected type, then it will pick one based on the least total cost for the prefix. For a tunneled next-hop, it means the LSP metric plus the cost of the LSP endpoint to the destination of the prefix.
  7. If there is more than one LFA next-hop within the selected type (ecmp-case) in the first set, it will select the first direct next-hop from the remaining set. This is not a deterministic selection and will vary following each SPF calculation.
  8. If there is more than one LFA next-hop within the selected type (ecmp-case) in the second set, it will pick the tunneled next-hop with the lowest cost from the endpoint of the LSP to the destination prefix. If there remains more than one, it will pick the tunneled next-hop with the lowest tunnel-id.

### 5.10.3.2 Loop-Free Alternate Calculation for Inter-Area/inter-Level Prefixes

When SPF resolves OSPF inter-area prefixes or IS-IS inter-level prefixes, it will compute an LFA backup next-hop to the same exit area/border router as used by the primary next-hop.

### 5.10.3.3 Loop-Free Alternate Shortest Path First (LFA SPF) Policies

An LFA SPF policy allows the user to apply specific criteria, such as admin group and SRLG constraints, to the selection of a LFA backup next-hop for a subset of prefixes that resolve to a specific primary next-hop. See more details in the section titled “*Loop-Free Alternate Shortest Path First (LFA SPF) Policies*” in the *Routing Protocols Guide*.

## 5.11 mLDP Fast Upstream Switchover

mLDP Fast Upstream Switchover allows a downstream LSR of a multicast LDP (mLDP) FEC to perform a fast switchover and source the traffic from another upstream LSR while IGP is converging due to a failure of the primary next-hop of the Point-to-Multipoint (P2MP) FEC. In a sense, it provides an upstream Fast-Reroute (FRR) capability for the mLDP packets. It does it at the expense of traffic duplication from two different upstream nodes into the node that performs the fast upstream switchover.


When this feature is enabled and LDP is resolving an mLDP FEC received from a downstream LSR, it checks if an Equal-Cost Multi-Path (ECMP) next-hop or a Loop-Free Alternate (LFA) next-hop exist to the root LSR node. 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. LDP then sends the corresponding labels to the upstream LSR nodes. In normal operation, the primary ILM accepts packets while the backup ILM drops them. If the node detects that the interface or the upstream LSR of the primary ILM is down, the backup ILM will then start accepting packets.

In order to make use of the ECMP next-hop, the user must configure the ECMP value in the system to at least two (2). In order to make use of the LFA next-hop, the user must enable LFA and IP FRR options under the IGP instance.

## 5.12 LDP FEC to BGP Label Route Stitching

The stitching of an LDP FEC to a BGP labeled route allows the LDP capable PE devices to offer services to PE routers in other areas or domains without the need to support BGP labeled routes.

This feature is used in a large network to provide services across multiple areas or autonomous systems. [Figure 63](#) shows a network with a core area and regional areas.

**Figure 63 Application of LDP to BGP FEC Stitching**

al\_0217

Specific /32 routes in a regional area are not redistributed into the core area. Therefore, only nodes within a regional area and the ABR nodes in the same area exchange LDP FECs. A PE router, for example, PE21, in a regional area learns the reachability of PE routers in other regional areas by way of RFC 3107 BGP labeled routes redistributed by the remote ABR nodes by way of the core area. The remote ABR then sets the next-hop self on the labeled routes before re-distributing them into the core area. The local ABR for PE2, for example, ABR3 may or may not set next-hop self when it re-distributes these labeled BGP routes from the core area to the local regional area.

When forwarding a service packet to the remote PE, PE21 inserts a VC label, the BGP route label to reach the remote PE, and an LDP label to reach either ABR3, if ABR3 sets next-hop self, or ABR1.

In the same network, an MPLS capable DSLAM also act as PE router for VLL services and will need to establish a PW to a PE in a different regional area by way of router PE21, acting now as an LSR. To achieve that, PE21 is required to perform the following operations:

- Translate the LDP FEC it learned from the DSLAM into a BGP labeled route and re-distribute it by way of iBGP within its area. This is in addition to redistributing the FEC to its LDP neighbors in the same area.
- Translate the BGP labeled routes it learns through iBGP into an LDP FEC and re-distribute it to its LDP neighbors in the same area. In the application in [Figure 63](#), the DSLAM requests the LDP FEC of the remote PE router using LDP Downstream on Demand (DoD).

- When a packet is received from the DSLAM, PE21 swaps the LDP label into a BGP label and pushes the LDP label to reach ABR3 or ABR1. When a packet is received from ABR3, the top label is removed and the BGP label is swapped for the LDP label corresponding to the DSLAM FEC.

## 5.12.1 Configuration

The user enables the stitching of routes between the LDP and BGP by configuring separately tunnel table route export policies in both protocols and enabling the advertising of RFC 3107 formatted labeled routes for prefixes learned from LDP FECs.

The route export policy in BGP instructs BGP to listen to LDP route entries in the CPM tunnel table. If a /32 LDP FEC prefix matches an entry in the export policy, BGP originates a BGP labeled route, stitches it to the LDP FEC, and re-distributes the BGP labeled route to its iBGP neighbors.

The user adds LDP FEC prefixes with the statement 'from protocol ldp' in the configuration of the existing BGP export policy at the global level, the peer-group level, or at the peer level using the commands:

- **config>router>bgp>export** *policy-name*
- **config>router>bgp>group>export** *policy-name*
- **config>router>bgp>group>neighbor>export** *policy-name*

To indicate to BGP to evaluate the entries with the 'from protocol ldp' statement in the export policy when applied to a specific BGP neighbor, use commands:

- **config>router>bgp>group>neighbor>family label-ipv4**
- **config>router>bgp>group>neighbor>advertise-ldp-prefix**

Without this, only core IPv4 routes learned from RTM are advertised as BGP labeled routes to this neighbor. And the stitching of LDP FEC to the BGP labeled route is not performed for this neighbor even if the same prefix was learned from LDP.

The tunnel table route export policy in LDP instructs LDP to listen to BGP route entries in the CPM Tunnel Table. If a /32 BGP labeled route matches a prefix entry in the export policy, LDP originates an LDP FEC for the prefix, stitches it to the BGP labeled route, and re-distributes the LDP FEC its iBGP neighbors.

The user adds BGP labeled route prefixes with the statement 'from protocol bgp' in the configuration of a new LDP tunnel table export policy using the command:

**config>router>ldp>export-tunnel-table** *policy-name*.

The 'from protocol' statement has an effect only when the protocol value is **ldp**. Policy entries with protocol values of **rsvp**, **bgp**, or any value other than **ldp** are ignored at the time the policy is applied to LDP.

## 5.12.2 Detailed LDP FEC Resolution

When an LSR receives a FEC-label binding from an LDP neighbor for a given specific FEC1 element, the following procedures are performed.

1. LDP installs the FEC if:
  - It was able to perform a successful exact match or a longest match, if aggregate-prefix-match option is enabled in LDP, of the FEC /32 prefix with a prefix entry in the routing table.
  - The advertising LDP neighbor is the next-hop to reach the FEC prefix.
2. When such a FEC-label binding has been installed in the LDP FIB, LDP will perform the following:
  - Program a push and a swap NHLFE entries in the egress data path to forward packets to FEC1.
  - Program the CPM tunnel table with a tunnel entry for the NHLFE.
  - Advertise a new FEC-label binding for FEC1 to all its LDP neighbors according to the global and per-peer LDP prefix export policies.
  - Install the ILM entry pointing to the swap NHLFE.
3. When BGP learns the LDP FEC by way of the CPM tunnel table and the FEC prefix exists in the BGP route export policy, it will perform the following:
  - Originate a labeled BGP route for the same prefix with this node as the next-hop and advertise it by way of iBGP to its BGP neighbors, for example, the local ABR/ASBR nodes, which have the **advertise-ldp-prefix** enabled.
  - Install the ILM entry pointing to the swap NHLFE programmed by LDP.

## 5.12.3 Detailed BGP Labeled Route Resolution

When an LSR receives a BGP labeled route by way of iBGP for a given specific /32 prefix, the following procedures are performed.

1. BGP resolves and installs the route in BGP if:
  - There exists an LDP LSP to the BGP neighbor, for example, the ABR or ASBR, which advertised it and which is the next-hop of the BGP labeled route.

2. Once the BGP route is installed, BGP programs the following:
  - Push NHLFE in the egress data path to forward packets to this BGP labeled route.
  - The CPM tunnel table with a tunnel entry for the NHLFE.
3. When LDP learns the BGP labeled route by way of the CPM tunnel table and the prefix exists in the new LDP tunnel table route export policy, it performs the following:
  - Advertise a new LDP FEC-label binding for the same prefix to its LDP neighbors according the global and per-peer LDP export prefix policies. If LDP already advertised a FEC for the same /32 prefix after receiving it from an LDP neighbor then no action is required. For LDP neighbors that negotiated LDP Downstream on Demand (DoD), the FEC is advertised only when this node receives a Label Request message for this FEC from its neighbor.
  - Install the ILM entry pointing the BGP NHLFE if a new LDP FEC-label binding is advertised. If an ILM entry exists and points to an LDP NHLFE for the same prefix then no update to ILM entry is performed. The LDP route has always preference over the BGP labeled route.

## 5.12.4 Data Plane Forwarding

When a packet is received from an LDP neighbor, the LSR swaps the LDP label into a BGP label and pushes the LDP label to reach the BGP neighbor, for example, ABR/ASBR, which advertised the BGP labeled route with itself as the next-hop.

When a packet is received from a BGP neighbor such as an ABR/ASBR, the top label is removed and the BGP label is swapped for the LDP label to reach the next-hop for the prefix.

## 5.13 LDP-SR Stitching for IPv4 prefixes (IS-IS)

This feature enables stitching between an LDP FEC and an SR node-SID route for the same IPv4 /32prefix.

## 5.13.1 LDP-SR Stitching Configuration

The user enables the stitching between an LDP FEC and an SR node-SID route for the same prefix by configuring the export of SR (LDP) tunnels from the CPM Tunnel Table Manager (TTM) into LDP (IGP).

In the LDP-to-SR data path direction, the existing tunnel table route export policy in LDP, which was introduced for LDP-BGP stitching, is enhanced to include support for exporting SR tunnels from the TTM to LDP. The user adds the new statement **from protocol isis** [**instance** *instance-id*] to the LDP tunnel table export policy:

**CLI Syntax:**    `config>router>ldp>export-tunnel-table policy-name`

The user can restrict the export to LDP of SR tunnels from a specific prefix list. The user can also restrict the export to a specific IGP instance by optionally specifying the instance ID in the **from** statement.

The **from protocol** statement has an effect only when the protocol value is **isis** or **bgp**.

Policy entries with any other protocol value are ignored at the time the policy is applied. If the user configures multiple **from** statements in the same policy or does not include the **from** statement but adds a default action of **accept**, then LDP will follow the TTM selection rules as described in the *Segment Routing Tunnel Management* section of the *Unicast Routing Protocol Guide* to select a tunnel to stitch the LDP ILM to:

- LDP selects the tunnel from the lowest TTM preference protocol.
- If IS-IS and BGP protocols have the same preference, then LDP selects the protocol using the default TTM protocol preference.
- Within the same IGP protocol, LDP selects the lowest instance ID.

When this policy is enabled in LDP, LDP listens to SR tunnel entries in the TTM. Whenever an LDP FEC primary next-hop cannot be resolved using an RTM route and a SR tunnel of type SR-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 also originates an FEC for the prefix and re-distributes it to its LDP and T-LDP peers. The latter allows an LDP FEC that is tunneled over a RSVP-TE LSP to have its ILM stitched to an SR tunnel endpoint. When a LDP FEC is stitched to a SR tunnel, packets forwarded will benefit from the protection of the LFA/remote LFA backup next-hop of the SR tunnel.

When resolving a FEC, LDP will prefer resolution in RTM over that in TTM when both resolutions are possible. In other words, the swapping of the LDP ILM to a LDP NHLFE is preferred over stitching it 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. Refer to the *Segment Routing Mapping Server* section of the *Unicast Routing Protocols Guide* for more information. Thus, a tunnel table export policy is not required. Instead, the user enables exporting to an IGP instance the LDP tunnels for FEC prefixes advertised by the mapping server using the following command:

**CLI Syntax:** `config>router>isis>segment-routing>export-tunnel-table  
ldp`

When this command is enabled in the segment-routing context of an IGP instance, IGP listens to LDP tunnel entries in the TTM. Whenever a /32 LDP tunnel destination matches a prefix for which IGP received a prefix-SID sub-TLV from a mapping server, it 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 either link LDP or T-LDP. In the latter, the stitching is performed to an LDP-over-RSVP tunnel and only supported when the **ldp-over-rsvp** option is enabled in IGP. It is not supported when the **rsvp-shortcut** option is enabled. When an SR tunnel is stitched to an LDP FEC, packets forwarded will benefit from the FRR protection of the LFA backup next-hop of the LDP FEC.


When resolving a node SID, IGP will prefer resolution of prefix SID received in an IP Reach TLV over a prefix SID received via the mapping server. In other words, the swapping of the SR ILM to a SR NHLFE is preferred over stitching it to a LDP tunnel endpoint. Refer to the Segment Routing Mapping Server Prefix SID Resolution section of the Unicast Routing Protocols Guide for more information about prefix SID resolution.

It is recommended to enable the **bfd-enable** option on the interfaces in both LDP and IGP instance contexts to speed up the failure detection and the activation of the LFA/remote-LFA backup next-hop in either direction. This is particularly true if the injected failure is a remote failure.

This feature is limited to IPv4 /32 prefixes in both LDP and SR.

## 5.13.2 Stitching in the LDP-to-SR Direction

The stitching in data-plane from the LDP-to-SR direction is based on the LDP module monitoring the TTM for a SR tunnel of a prefix matching an entry in the LDP TTM export policy.

**Figure 64**     **Stitching in the LDP-to-SR Direction**

0970

With reference to [Figure 64](#), the following procedure is performed by the boundary router R1 to effect stitching:

- Step 1.** Router R1 is at the boundary between an SR domain and an LDP domain and is configured to stitch between SR and LDP.
- Step 2.** Link R1-R2 is LDP-enabled, but router R2 does not support SR (or SR is disabled).
- Step 3.** Router R1 receives a prefix-SID sub-TLV in an IS-IS IP reachability TLV originated by router Ry for prefix Y.
- Step 4.** R1 resolves it and programs an NHLFE on the link towards the next-hop in the SR domain. It programs an SR ILM and points it to this NHLFE.
- Step 5.** Because R1 is programmed to stitch LDP to SR, the LDP in R1 discovers in TTM the SR tunnel to Y. LDP programs a 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.
- Step 6.** R1 advertises the LDP FEC for the prefix Y to all its LDP peers. R2 is now able to install a LDP tunnel towards Ry.
- Step 7.** If R1 found multiple SR tunnels to destination prefix Y, it should use the TTM tunnel selection rules to pick the SR tunnel. The rules follow the following steps:
  - i. R1 selects the tunnel from the lowest preference IGP protocol.
  - ii. Select the protocol using the default TTM protocol preference.
  - iii. Within the same IGP protocol, R1 uses the lowest instance ID to select the tunnel.

- Step 8.** If the user configured in the same LDP tunnel table export policy concurrently **from protocol isis** and **from protocol bgp**, or did not include the from statement but added a default action of accept, R1 will select the tunnel to destination prefix Y to stitch the LDP ILM to using the TTM tunnel selection rules:
- R1 selects the tunnel from the lowest preference protocol.
  - If IS-IS and BGP protocols have the same preference, then R1 selects the protocol using the default TTM protocol preference.
  - Within the same IGP protocol, R1 uses the lowest instance ID to select the tunnel.



**Note:** If R1 has already resolved a LDP FEC for prefix Y, it would have had an ILM for it, but this ILM is not be updated to point towards the SR tunnel. This is because LDP resolves in RTM first before going to TTM and thus prefers the LDP tunnel over the SR tunnel. Similarly, if an LDP FEC is received subsequently to programming the stitching, the LDP ILM will be updated to point to the LDP NHLFE because LDP will be able to resolve the LDP FEC in RTM.

- Step 9.** The user enables SR in R2. R2 resolves the prefix SID for Y and installs the SR ILM and the SR NHLFE. R2 is now able of forwarding packets over the SR tunnel to router Ry. Nothing happens in R1 because the SR ILM is already programmed.
- Step 10.** The user disables LDP on the interface R1-R2 (both directions) and the LDP FEC ILM and NHLFE are removed in R1. The same occurs in R2 which can then only forward using the SR tunnel towards Ry.

### 5.13.3 Stitching in the SR-to-LDP Direction

The stitching in data-plane from the SR-to-LDP direction is based on the IGP monitoring the TTM for a LDP tunnel of a prefix matching an entry in the SR TTM export policy.

With reference to [Figure 64](#), the following procedure is performed by the boundary router R1 to effect stitching:

- Step 1.** Router R1 is at the boundary between a SR domain and a 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).
- Step 2.** R1 receives an LDP FEC for prefix X owned by router Rx further down in the LDP domain.

RTM in R1 shows that the interface to R2 is the next-hop for prefix X.

- Step 3.** LDP in R1 resolves this FEC in RTM and creates an LDP ILM for it with, for example, ingress label L1, and points it to an LDP NHLFE towards R2 with egress label L2.
- Step 4.** Later on, R1 receives a prefix-SID sub-TLV from the mapping server R5 for prefix X.
- Step 5.** IGP in R1 is resolving in its routing table the next-hop of prefix X to the interface to R2. R1 knows that R2 did not advertise support of Segment Routing and, thus, SID resolution for prefix X in routing table fails.
- Step 6.** IGP in R1 attempts to resolve prefix SID of X in TTM because it is configured to stitch SR-to-LDP. R1 finds a LDP tunnel to X in TTM, instructs the SR module to program a SR ILM with ingress label L3, and points it to the LDP tunnel endpoint, thus stitching ingress L3 label to egress L2 label.

**Note:**



- Here, two ILMs, the LDP and SR, are pointing to the same LDP tunnel one via NHLFE and one via tunnel endpoint.
- No SR tunnel to destination X should be programmed in TTM following this resolution step.
- A trap will be generated for prefix SID resolution failure only after IGP fails to complete Step 5 and Step 6. The existing trap for prefix SID resolution failure is enhanced to state whether the prefix SID which failed resolution was part of mapping server TLV or a IP reachability TLV (ISIS).

**Step 7.** The user enables segment routing on R2.

**Step 8.** IGP in R1 discovers that R2 supports SR via the SR capability.

Because R1 still has a prefix-SID for X from the mapping server R5, it maintains the stitching of the SR ILM for X to the LDP FEC unchanged.

**Step 9.** The operator disables the LDP interface between R1 and R2 (both directions) and the LDP FEC ILM and NHLFE for prefix X are removed in R1.

**Step 10.** This triggers the re-evaluation of the SIDs. R1 first attempts the resolution in routing table and since the next-hop for X now supports SR, IGP instructs the SR module to program a NHLFE for prefix-SID of X with egress label L4 and outgoing interface to R2. R1 installs a SR tunnel in TTM for destination 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.

**Step 11.** Much later on, the router that owns prefix X, Rx, was upgraded to support SR. R1 now receives a prefix-SID sub-TLV in a ISIS IP reachability TLV originated by Rx for prefix X. The SID information may or may not be the same as the one received from the mapping server R5. In this case, IGP in R1 will prefer the prefix-SID originated by Rx and will thus update the SR ILM and NHLFE with appropriate labels.

**Step 12.** Finally, the operator cleans up the mapping server and removes the mapping entry for prefix X, which then gets withdrawn by IS-IS.

## 5.14 LDP FRR Remote LFA Backup using SR Tunnel for IPv4 Prefixes (IS-IS)

The user enables the use of SR tunnel as a remote LFA backup tunnel next-hop by an LDP FEC via the following CLI command:

**CLI Syntax:** `config>router>ldp>fast-reroute [backup-sr-tunnel]`

As a pre-requisite, the user must enable the stitching of LDP and SR in the LDP-to-SR direction as explained in [5.13.1](#). That is because the LSR must perform the stitching of the LDP ILM to SR tunnel when the primary LDP next-hop of the FEC fails. Thus, LDP must listen to SR tunnels programmed by the IGP in TTM, but the mapping server feature is not required.

Assume the **backup-sr-tunnel** option is enabled in LDP and the **{loopfree-alternate remote-lfa}** option is enabled in the IGP instance, and that LDP was able to resolve the primary next-hop of the LDP FEC in RTM. If the IGP LFA SPF does not find a regular LFA backup next-hop for a prefix of an LDP FEC, it will run the remote LFA algorithm. If IGP finds a remote LFA tunnel next-hop, LDP programs the primary next-hop of the FEC using an LDP NHLFE and programs the LFA backup next-hop using an LDP NHLFE pointing to the SR tunnel endpoint.



**Note:** The LDP packet is not “tunneled” over the SR tunnel. The LDP label is actually stitched to the segment routing label stack. LDP points both the LDP ILM and the LTN to the backup LDP NHLFE which itself uses the SR tunnel endpoint.

The behavior of the feature is similar to the LDP-to-SR stitching feature described in the [5.13LDP-SR Stitching for IPv4 prefixes \(IS-IS\)](#) section, except the behavior is augmented to allow the stitching of an LDP ILM/LTN to an SR tunnel also when the primary LDP next-hop of the FEC fails.

The following is the behavior of this feature:

- When LDP resolves a primary next-hop in RTM and a remote LFA backup next-hop using SR tunnel in TTM, LDP programs a primary LDP NHLFE as usual and a backup LDP NHLFE pointing to the SR tunnel, which has the remote LFA backup for the same prefix.
- If the LDP FEC primary next-hop failed and LDP has pre-programmed a remote LFA next-hop with an LDP backup NHLFE pointing to the SR tunnel, the LDP ILM/LTN switches to it.



**Note:** If, for some reason, the failure impacted only the LDP tunnel primary next-hop but not the SR tunnel primary next-hop, the LDP backup NHLFE will effectively point to the primary next-hop of the SR tunnel and traffic of the LDP ILM/LTN will follow this path instead of the remote LFA next-hop of the SR tunnel until the latter is activated.

- If the LDP FEC primary next-hop becomes unresolved in RTM, LDP switches the resolution to a SR tunnel in TTM, if one exists, as per the LDP-to-SR stitching procedures described in [5.13.2](#).
- If both the LDP primary next-hop and a regular LFA next-hop become resolved in RTM, the LDP FEC programs the primary and backup NHLFEs as usual.
- It is recommended to enable the **bfd-enable** option on the interfaces in both LDP and IGP instance contexts to speed up the failure detection and the activation of the LFA/remote-LFA backup next-hop in either direction.

## 5.15 Automatic Creation of a Targeted Hello Adjacency and LDP Session

This feature enables the automatic creation of a targeted Hello adjacency and LDP session to a discovered peer.

### 5.15.1 Feature Configuration

The user first creates a targeted LDP session peer parameter template:

```
config>router>ldp>targ-session>peer-template template-name
```

Inside the template the user configures the common T-LDP session parameters or options shared by all peers using this template. These are the following:

**bfd-enable**, **hello**, **hello-reduction**, **keepalive**, **local-lsr-id**, and **tunneling**.

The tunneling option does not support adding explicit RSVP LSP names. LDP will select RSVP LSP for an endpoint in LDP-over-RSVP directly from the Tunnel Table Manager (TTM).

Then the user references the peer prefix list which is defined inside a policy statement defined in the global policy manager.

**config>router>ldp>targ-session>peer-template-map peer-template *template-name* policy *peer-prefix-policy***

Each application of a targeted session template to a given prefix in the prefix list will result in the establishment of a targeted Hello adjacency to an LDP peer using the template parameters as long as the prefix corresponds to a router-id for a node in the TE database. The targeted Hello adjacency will either trigger a new LDP session or will be associated with an existing LDP session to that peer.

Up to five (5) peer prefix policies can be associated with a single peer template at all times. Also, the user can associate multiple templates with the same or different peer prefix policies. Thus multiple templates can match with a given peer prefix. In all cases, the targeted session parameters applied to a given peer prefix are taken from the first created template by the user. This provides a more deterministic behavior regardless of the order in which the templates are associated with the prefix policies.

Each time the user executes the above command, with the same or different prefix policy associations, or the user changes a prefix policy associated with a targeted peer template, the system re-evaluates the prefix policy. The outcome of the re-evaluation will tell LDP if an existing targeted Hello adjacency needs to be torn down or if an existing targeted Hello adjacency needs to have its parameters updated on the fly.

If a /32 prefix is added to (removed from) or if a prefix range is expanded (shrunk) in a prefix list associated with a targeted peer template, the same prefix policy re-evaluation described above is performed.

The template comes up in the **no shutdown** state and as such it takes effect immediately. Once a template is in use, the user can change any of the parameters on the fly without shutting down the template. In this case, all targeted Hello adjacencies are.

There is no overall chassis mode restrictions enforced with the auto-created T-LDP session feature. If the chassis-mode, network chassis-mode or IOM type requirements for an LDP feature are not met, the configuration of the corresponding command will not be allowed as in existing implementation.

## 5.15.2 Feature Behavior

Whether the prefix list contains one or more specific /32 addresses or a range of addresses, an external trigger is required to indicate to LDP to instantiate a targeted Hello adjacency to a node which address matches an entry in the prefix list. The objective of the feature is to provide an automatic creation of a T-LDP session to the same destination as an auto-created RSVP LSP to achieve automatic tunneling of LDP-over-RSVP. The external trigger is when the router with the matching address appears in the Traffic Engineering database. In the latter case, an external module monitoring the TE database for the peer prefixes provides the trigger to LDP. As a result of this, the user must enable the **traffic-engineering** option in ISIS or OSPF.

Each mapping of a targeted session peer parameter template to a policy prefix which exists in the TE database will result in LDP establishing a targeted Hello adjacency to this peer address using the targeted session parameters configured in the template. This Hello adjacency will then either get associated with an LDP session to the peer if one exists or it will trigger the establishment of a new targeted LDP session to the peer.

The SR OS supports multiple ways of establishing a targeted Hello adjacency to a peer LSR:

- User configuration of the peer with the targeted session parameters inherited from the **config>router>ldp>targ-session>ipv4** in the top level context or explicitly configured for this peer in the **config>router>ldp>targ-session>peer** context and which overrides the top level parameters shared by all targeted peers. Let us refer to the top level configuration context as the global context. Some parameters only exist in the global context; their value will always be inherited by all targeted peers regardless of which event triggered it.
- User configuration of an SDP of any type to a peer with the **signaling tldp** option enabled (default configuration). In this case the targeted session parameter values are taken from the global context.
- User configuration of a (FEC 129) PW template binding in a BGP-VPLS service. In this case the targeted session parameter values are taken from the global context.
- User configuration of a (FEC 129 type II) PW template binding in a VLL service (dynamic multi-segment PW). In this case the target session parameter values are taken from the global context
- This Release 11.0.R4 user configuration of a mapping of a targeted session peer parameter template to a prefix policy when the peer address exists in the TE database. In this case, the targeted session parameter values are taken from the template.



- Features using an LDP LSP, which itself is tunneled over an RSVP LSP (LDP-over-RSVP), as a shortcut do not trigger automatically the creation of the targeted Hello adjacency and LDP session to the destination of the RSVP LSP. The user must configure manually the peer parameters or configure a mapping of a targeted session peer parameter template to a prefix policy. These features are:
  - BGP shortcut (**next-hop-resolution shortcut-tunnel** option in BGP),
  - IGP shortcut (**rsvp-shortcut** option in IGP),
  - LDP shortcut for IGP routes (**ldp-shortcut** option in router level),
  - static route LDP shortcut (**ldp** option in a static route),
  - VPRN service (**auto-bind-tunnel ldp** option), and

Since the above triggering events can occur simultaneously or in any arbitrary order, the LDP code implements a priority handling mechanism in order to decide which event overrides the active targeted session parameters. The overriding trigger will become the owner of the targeted adjacency to a given peer and will be shown in **show router ldp targ-peer**.

[Table 39](#) summarizes the triggering events and the associated priority.

**Table 39 Targeted LDP Adjacency Triggering Events and Priority**

| Triggering Event                                                                                | Automatic Creation of Targeted Hello Adjacency | Active Targeted Adjacency Parameter Override Priority |
|-------------------------------------------------------------------------------------------------|------------------------------------------------|-------------------------------------------------------|
| Manual configuration of peer parameters (creator=manual)                                        | Yes                                            | 1                                                     |
| Mapping of targeted session template to prefix policy (creator=template)                        | Yes                                            | 2                                                     |
| Manual configuration of SDP with <b>signaling tldp</b> option enabled (creator=service manager) | Yes                                            | 3                                                     |
| PW template binding in BGP-AD VPLS (creator=service manager)                                    | Yes                                            | 3                                                     |
| PW template binding in FEC 129 VLL (creator=service manager)                                    | Yes                                            | 3                                                     |
| LDP-over-RSVP as a BGP/IGP/LDP/Static shortcut                                                  | No                                             | N/A                                                   |
| LDP-over-RSVP in VPRN auto-bind                                                                 | No                                             | N/A                                                   |
| LDP-over-RSVP in BGP Label Route resolution                                                     | No                                             | N/A                                                   |

Any parameter value change to an active targeted Hello adjacency caused by any of the above triggering events is performed by having LDP immediately send a Hello message with the new parameters to the peer without waiting for the next scheduled time for the Hello message. This allows the peer to adjust its local state machine immediately and maintains both the Hello adjacency and the LDP session in UP state. The only exceptions are the following:

- The triggering event caused a change to the **local-lsr-id** parameter value. In this case, the Hello adjacency is brought down which will also cause the LDP session to be brought down if this is the last Hello adjacency associated with the session. A new Hello adjacency and LDP session will then get established to the peer using the new value of the local LSR ID.
- The triggering event caused the targeted peer **shutdown** option to be enabled. In this case, the Hello adjacency is brought down which will also cause the LDP session to be brought down if this is the last Hello adjacency associated with the session.


Finally, the value of any LDP parameter which is specific to the LDP/TCP session to a peer is inherited from the **config>router>ldp>session-params>peer** context. This includes MD5 authentication, LDP prefix per-peer policies, label distribution mode (DU or DOD), etc.

## 5.16 Multicast P2MP LDP for GRT

The P2MP LDP LSP setup is initiated by each leaf node of multicast tree. A leaf PE node learns to initiate a multicast tree setup from client application and sends a label map upstream towards the root node of the multicast tree. On propagation of label map, intermediate nodes that are common on path for multiple leaf nodes become branch nodes of the tree.

[Figure 65](#) illustrates wholesale video distribution over P2MP LDP LSP. Static IGMP entries on edge are bound to P2MP LDP LSP tunnel-interface for multicast video traffic distribution.

**Figure 65 Video Distribution using P2MP LDP**



## 5.17 LDP P2MP Support

### 5.17.1 LDP P2MP Configuration

A node running LDP also supports P2MP LSP setup using LDP. By default, it would advertise the capability to a peer node using P2MP capability TLV in LDP initialization message.

This configuration option per interface is provided to restrict/allow the use of interface in LDP multicast traffic forwarding towards a downstream node. Interface configuration option does not restrict/allow exchange of P2MP FEC by way of established session to the peer on an interface, but it would only restrict/allow use of next-hops over the interface.

## **5.17.2 LDP P2MP Protocol**

Only a single generic identifier range is defined for signaling multipoint tree for all client applications. Implementation on the 7750 SR or 7950 XRS reserves the range (1..8292) of generic LSP P2MP-ID on root node for static P2MP LSP.

## **5.17.3 Make Before Break (MBB)**

When a transit or leaf node detects that the upstream node towards the root node of multicast tree has changed, it follows graceful procedure that allows make-before-break transition to the new upstream node. Make-before-break support is optional. If the new upstream node does not support MBB procedures then the downstream node waits for the configured timer before switching over to the new upstream node.

## **5.17.4 ECMP Support**

If multiple ECMP paths exist between two adjacent nodes on the then the upstream node of the multicast receiver programs all entries in forwarding plane. Only one entry is active based on ECMP hashing algorithm.

## 5.17.5 Inter-AS Non-segmented mLDP

This feature allows multicast services to use segmented protocols and span them over multiple autonomous systems (ASs), as done in unicast services. As IP VPN or GRT services span multiple IGP areas or multiple ASs, either due to 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. SR OS allows for similar provision of multicast services and for spanning these services over multiple IGP areas or multiple ASs.

For multicast VPN (MVPN), SR OS previously supported inter-AS option A/B/C for Rosen MVPN; however, when MPLS was used, only option A was supported for Next Generation Multicast VPN (NG-MVPN) or d-mLDP signaling. mLDP now supports non-segmented mLDP trees for inter-AS solutions, applicable for multicast services in the GRT (Global Routing Table) where they need to traverse mLDP point-to-multipoint tunnels as well as NG-MVPN services.

### 5.17.5.1 In-band Signaling with Non-segmented mLDP Trees in GRT

mLDP can be used to transport multicast in GRT. This functionality can be configured under IGMP or, in MVPN, under inclusive or selective provider tunnels. In either case, for mLDP Label Switched Paths (LSP) to be generated, a multicast request from the leaf node is required to force mLDP to generate a DU (downstream unsolicited) FEC toward the root to build the P2MP LSPs.


For inter-AS solutions, the root might not be in the leaf's RTM or, if it is present, it is installed via BGP with ASBRs acting as the leaf's local AS root. As such, the leaf's local AS intermediate routers might not know the path to the root.

Summarily, control protocols used for constructing P2MP LSPs contain a field that identifies the address of a root node. Intermediate nodes are expected to be able to look up that address in their routing tables; however, this is not possible if the route to the root node is a BGP route and the intermediate nodes are part of a BGP-free core (for example, if they use IGP).

To enable an mLDP LSP to be constructed through a BGP-free segment, the root node address is temporarily replaced by an address that is known to the intermediate nodes and is on the path to the true root node. For example, [Figure 66](#) shows the procedure when the PE-2 (leaf) receives the route for root via ASBR3. This route looks like the following root next-hop ASBR-3. The leaf, in this case, generates an

LDP FEC which has an opaque value, and has the root address set as ASBR-3. This opaque value has additional information needed to reach the root from ASBR-3. As a result, the SR core AS3 only needs to be able to resolve the local AS ASBR-3 for the LDP FEC. The ASBR-3 uses the LDP FEC opaque value to find the path to the root.

**Figure 66 Inter-AS Option C**



As non-segmented d-mLDP requires end-to-end mLDP signaling, the ASBRs support both mLDP and BGP signaling between them (hence the term “non-segmented”, as the mLDP signaling needs to be from the leaf to the root node).

### 5.17.5.2 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 via 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 P2MP FEC element, encoded as specified in RFC 6513, with a type field, a length field, and a value field of its own.

For example, RFC 6826 section 3 defines the Transit IPv4 opaque for P2MP LDP FEC, where the leaf in the local AS wants to establish an LSP to the root for P2MP LSP. [Figure 67](#) shows this FEC representation.

**Figure 67 mLDP FEC for Single AS with Transit IPv4 Opaque**






Figure 68 shows an inter-AS FEC with recursive opaque based on RFC 6512.

**Figure 68 mLDP FEC for Inter-AS with Recursive Opaque Value**



As shown in Figure 68, the root “100.0.0.21” is an ASBR and the opaque value contains the original mLDP FEC. As such, in the leaf’s AS where the actual root “100.0.0.14” is not known, the LDP FEC can be routed via the local root of ASBR. When the FEC arrives at an ASBR that co-locates in the same AS as the actual root, an LDP FEC with transit IPv4 opaque is generated. The end-to-end picture for inter-AS mLDP for non-VPN multicast shown in Figure 69.

**Figure 69 Non-VPN mLDP with Recursive Opaque for Inter-AS**

1027

As shown in [Figure 68](#), the leaf is in AS3s where the AS3 intermediate nodes do not have the ROOT-1 in their RTM. The leaf has the S1 installed in the RTM via BGP. All ASBRs are acting as next-hop-self in the BGP domain. The leaf resolving the S1 via BGP will generate an mLDP FEC with recursive opaque, represented as:

**Leaf FEC: <Root=ASBR-3, opaque-value=<Root=Root-1, <opaque-value = S1,G1>>>**

This FEC will be routed through the AS3 Core to ASBR-3.



**Note:** AS3 intermediate nodes do not have ROOT-1 in their RTM; that is, are not BGP-capable.

At ASBR-3 the FEC will be changed to:

**ASBR-3 FEC: <Root=ASBR-1, opaque-value=<Root=Root-1, <opaque-value = S1,G1>>>**

This FEC will be routed from ASBR-3 to ASBR-1. ASBR-1 is co-located in the same AS as ROOT-1 (that is, both ROOT-1 and ASBR-1 are in AS1). Therefore, the 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. All d-mLDP opaque types can be used in a FEC with a recursive opaque value.

### 5.17.5.3 Supported Recursive Opaque Values

A recursive FEC is built via opaque type Recursive Opaque Basic and VPN (opaque values 7 and 8 respectively). All SR non-recursive opaque values can be recursively embedded into a Recursive Opaque.

[Table 40](#) displays all supported opaque values in SR OS.


**Table 40** Recursive Opaque Types

| Opaque Type | Opaque Name                   | RFC      | SR OS Use                        |
|-------------|-------------------------------|----------|----------------------------------|
| 1           | Basic Type                    | RFC 6388 | VPRN Local AS                    |
| 3           | Transit IPv4                  | RFC 6826 | IPv4 multicast over mLDP in GRT  |
| 4           | Transit IPv6                  | RFC 6826 | IPv6 multicast over mLDP in GRT  |
| 7           | Recursive Opaque (Basic Type) | RFC 6512 | Inter-AS Option C MVPN over mLDP |
| 8           | Recursive Opaque (VPN Type)   | RFC 6512 | Inter-AS Option B MVPN over mLDP |

### 5.17.5.4 Redundancy and Resiliency

For mLDP, MoFRR is supported with the IGP domain; for example, ASBRs that are not directly connected. MoFRR is not supported between directly connected ASBRs, such as ASBRs that are using eBGP without IGP.

Figure 70 ASBRs Using eBGP Without IGP



5.17.5.5 ASBR Physical Connection


Non-segmented mLDP functions with ASBRs directly connected or connected via an IGP domain, as shown in [Figure 70](#).

5.17.5.6 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 via the vanilla IP to the source. Similarly, for a **p2mp-lsp-ping**, on the root, the echo request is transmitted via the mLDP P2MP tunnel to all leafs and the leafs use vanilla IP to respond to the root.

The echo request for mLDP is generated carrying a root Target FEC Stack TLV, which is used to identify the multicast LDP LSP under test at the leaf. The Target FEC Stack TLV must carry an mLDP P2MP FEC Stack Sub-TLV from RFC 6388 or RFC 6512.


Figure 71 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.

For inter-AS Option B where the root is not present in the leaf RTM, the echo reply cannot be forwarded via the GRT to the root. To solve this problem, for inter-AS Option B, the SR OS uses VPRN unicast routing to transmit the echo reply from the leaf to the root via VPRN.

**Figure 72 MVPN Inter-AS Option B OAM**



As shown in [Figure 72](#), the echo request for VPN recursive FEC is generated from the root node by executing the **p2mp-lsp-ping** with the **vpn-recursive-fec** option. When the echo request reaches the leaf, the leaf uses the sub-TLV within the echo request to identify the corresponding VPN via the FEC which includes the RD, the root, and the P2MP-ID.

After identifying the VPRN, the echo response is sent back via the VPRN and unicast routes. There should be a unicast route (for example, root 100.0.0.14, as shown in [Figure 72](#)) present in the leaf VPRN to allow the unicast routing of the echo reply back to the root via VPRN. To distribute this root from the root VPRN to all VPRN leafs, a loopback interface should be configured in the root VPRN and distributed to all leafs via MP-BGP unicast routes.

Notes:

1. For SR OS, all P2MP mLDP FEC types will respond to the **vpn-recursive-fec** echo request. Leafs in the local-AS and inter-AS Option C will respond to the recursive-FEC TLV echo request in addition to the leafs in the inter-AS Option B.

- a. For non inter-AS Option B where the root system IP is visible through the GRT, the echo reply will be sent via the GRT, that is, not via the VPRN.
2. This **vpn-recursive-fec** is a Nokia proprietary implementation, and therefore third-party routers will not recognize the recursive FEC and will not generate an echo respond.
  - a. The user can generate the **p2mp-lsp-ping** without the **vpn-recursive-fec** to discover non-Nokia routers in the local-AS and inter-AS Option C, but not the inter-AS Option B leafs.


**Table 41** OAM Functionality for Options B and C

| OAM Command (for mLDP)              | Leaf and Root in Same AS | Leaf and Root in Different AS (Option B) | Leaf and Root in Different AS (Option C) |
|-------------------------------------|--------------------------|------------------------------------------|------------------------------------------|
| p2mp-lsp-ping ldp                   | ✓                        | X                                        | ✓                                        |
| p2mp-lsp-ping ldp-ssm               | ✓                        | X                                        | ✓                                        |
| p2mp-lsp-ping ldp vpn-recursive-fec | ✓                        | ✓                                        | ✓                                        |
| p2mp-lsp-trace                      | X                        | X                                        | X                                        |

### 5.17.5.7 ECMP Support

In [Figure 73](#), the leaf discovers the ROOT-1 from all three ASBRs (ASBR-3, ASBR-4 and ASBR-5).

**Figure 73 ECMP Support**



1035

The leaf chooses which ASBR will be 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 configured ECMP (**config>router>ecmp**). For example, if the ECMP value is set to 2, only two of the ASBRs will be part of the hash algorithm selection.

2. After deciding 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 value in **config>router>ecmp**. 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 the ASBR 5 was chosen in [Figure 73](#), there are three paths between the leaf and ASBR-5. As such, a second ECMP decision is made to choose the path.
3. At ASBR-5, the process is repeated. For example, in [Figure 73](#), ASBR-5 will go through steps 1 and 2 to choose between ASBR-1 and ASBR-2, and 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 matter. If the LSR selection is done over a LAN interface and the Section 6 procedures are applied, the procedure described in [ECMP Hash Algorithm](#) should be applied to ensure that the same upstream LSR is elected among a 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.

#### 5.17.5.7.1 ECMP Hash Algorithm

The ECMP hash algorithm is based on RFC 6388 section 2.4.1.1.

1. The candidate upstream LSRs are numbered from lower to higher IP addresses.
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 identified in the FEC element after “Opaque Length”. The “Opaque Length” indicates the size of the opaque value used in this calculation.
3. The selected upstream LSR U is the LSR that has the number H.

#### 5.17.5.8 Dynamic mLDP and Static mLDP Co-existing on the Same Node

When creating a static mLDP tunnel, the user must configure the P2MP tunnel ID.

**Example:**

```
*A:SwSim2>config>router# tunnel-interface
no tunnel-interface ldp-p2mp p2mp-id sender sender-
address
tunnel-interface ldp-p2mp p2mp-id sender sender-
address [root-node]
```

This **p2mp-id** can coincide with a dynamic mLDP **p2mp-id** (dynamic mLDP is created by PIM automatically without configuration required). In scenarios where the node has a static mLDP and dynamic mLDP with same label and **p2mp-id**, there will be collisions and OAM misbehavior.

It is recommended not to use the static mLDP and dynamic mLDP on same node and, if it is done, to ensure that the **p2mp-id** is not the same between the two tunnel types.

---

## 5.18 Multicast LDP Fast Upstream Switchover

This feature allows a downstream LSR of a multicast LDP (mLDP) FEC to perform a fast switchover and source the traffic from another upstream LSR while IGP and LDP are converging due to a failure of the upstream LSR which is the primary next-hop of the root LSR for the P2MP FEC. In essence it provides an upstream Fast-Reroute (FRR) node-protection capability for the mLDP FEC packets. It does it at the expense of traffic duplication from two different upstream nodes into the node which performs the fast upstream switchover.

The detailed procedures for this feature are described in *draft-pdutta-mpls-mldp-up-redundancy*.

### 5.18.1 Feature Configuration

The user enables the mLDP fast upstream switchover feature by configuring the following option in CLI:

```
config>router>ldp>mcast-upstream-frr
```

When this command is enabled and LDP is resolving a mLDP FEC received from a downstream LSR, it checks if an ECMP next-hop or a LFA next-hop exist to the root LSR node. 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. LDP then sends the corresponding labels to both upstream LSR nodes. In 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 will then start accepting packets.

In order to make use of the ECMP next-hop, the user must configure the **ecmp** value in the system to at least two (2) using the following command:

```
config>router>ecmp
```

In order to make use of the LFA next-hop, the user must enable LFA using the following commands:

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

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

Enabling IP FRR or LDP FRR using the following commands is not strictly required since LDP only needs to know where the alternate next-hop to the root LSR is to be able to send the Label Mapping message to program the backup ILM at the initial signaling of the tree. Thus enabling the LFA option is sufficient. If however, unicast IP and LDP prefixes need to be protected, then these features and the mLDP fast upstream switchover can be enabled concurrently:

```
config>router>ip-fast-reroute
```

```
config>router>ldp>fast-reroute
```



**Caution:** The mLDP FRR fast switchover relies on the fast detection of loss of \*\*LDP session\*\* to the upstream peer to which the primary ILM label had been advertised. We strongly recommend that you perform 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 will bring down immediately the LDP session which will cause the IOM to activate the backup ILM.
2. If there is a concurrent TLDP adjacency to the same upstream LSR node, enable BFD on the T-LDP peer in addition to enabling it on the interface.
3. Enable 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 other reason than a failure of the interface or of the upstream LSR, routing and LDP will go out of sync. This means the backup ILM will remain activated until the next time SPF is rerun by IGP. By enabling IGP-LDP synchronization feature, the advertised link metric will be changed to max value as soon as the LDP session goes down. This in turn will trigger an SPF and LDP will likely download a new set of primary and backup ILMs.

## 5.18.2 Feature Behavior

This feature allows a downstream LSR to send a label binding to a couple of upstream LSR nodes but only accept traffic from the ILM on the interface to the primary next-hop of the root LSR for the P2MP FEC in normal operation, and accept traffic from the ILM on the interface to the backup next-hop under failure. Obviously, a candidate upstream LSR node must either be an ECMP next-hop or a Loop-Free Alternate (LFA) next-hop. This allows the downstream LSR to perform a fast switchover and source the traffic from another upstream LSR while IGP is converging due to a failure of the LDP session of the upstream peer which is the primary next-hop of the root LSR for the P2MP FEC. In a sense it provides an upstream Fast-Reroute (FRR) node-protection capability for the mLDP FEC packets.



**Figure 74 mLDP LSP with Backup Upstream LSR Nodes**



Upstream LSR **U** in [Figure 74](#) is the primary next-hop for the root LSR **R** of the P2MP FEC. This is also referred to as primary upstream LSR. Upstream LSR **U'** is an ECMP or LFA backup next-hop for the root LSR **R** of the same P2MP FEC. This is referred to as backup upstream LSR. 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 a backup ILM on the interface to LSR **U'**. The labels for the primary and backup ILMs must be different. LSR **Z** thus will attract traffic from both of them. However, LSR **Z** will block the ILM on the interface to LSR **U'** and will only accept traffic from the ILM on the interface to LSR **U**.

In case of a failure of the link to LSR **U** or of the LSR **U** itself causing the LDP session to LSR **U** to go down, LSR **Z** will detect it and reverse the ILM blocking state and will immediately start receiving traffic from LSR **U'** until IGP converges and provides a new primary next-hop, and ECMP or LFA backup next-hop, which may or may not be on the interface to LSR **U'**. At that point LSR **Z** will update the primary and backup ILMs in the data path.

The LDP uses the interface of either an ECMP next-hop or a LFA next-hop to the root LSR prefix, whichever is available, to program the backup ILM. ECMP next-hop and LFA next-hop are however mutually exclusive for a given prefix. IGP installs the ECMP next-hop in preference to an LFA next-hop for 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:

1. The candidate upstream LSRs are numbered from lower to higher 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 identified in the P2MP FEC Element right after 'Opaque Length' field. The 'Opaque Length' indicates the size of 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 ($H + 1 < \text{NUM\_ECMP}$) {
```

```
// 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 pickup the first.
```

```
 backup = 1;
```

```
}
```

In some topologies, it is possible that none of ECMP or LFA next-hop will be found. In this case, LDP programs the primary ILM only.

### 5.18.3 Uniform Failover from Primary to Backup ILM

When LDP programs the primary ILM record in the data path, it provides the IOM with the Protect-Group Identifier (PG-ID) associated with this ILM and which identifies which upstream LSR is protected.

In order for the system to perform a fast switchover to the backup ILM in the fast path, LDP applies to the primary ILM uniform FRR failover procedures similar in concept to the ones applied to an NHLFE in the existing implementation of LDP FRR for unicast FECs. There are however important differences to note. LDP associates a unique Protect Group ID (PG-ID) to all mLDP FECs which have their primary ILM on any LDP interface pointing **to the same upstream LSR**. This PG-ID is assigned per upstream LSR regardless of the number of LDP interfaces configured to this LSR. As such this PG-ID is different from the one associated with unicast FECs and which is assigned to each downstream LDP interface and next-hop. If however a failure

caused an interface to go down and also caused the LDP session to upstream peer to go down, both PG-IDs have their state updated in the IOM and thus the uniform FRR procedures will be triggered for both the unicast LDP FECs forwarding packets towards the upstream LSR and the mLDP FECs receiving packets from the same upstream LSR.

When the mLDP FEC is programmed in the data path, the primary and backup ILM record thus contain the PG-ID the FEC is associated with. The IOM also maintains a list of PG-IDs and a state bit which indicates if it is UP or DOWN. When the PG-ID state is UP the primary ILM for each mLDP FEC is open and will accept mLDP packets while the backup ILM is blocked and drops mLDP packets. LDP sends a PG-ID DOWN notification to IOM when it detects that the LDP session to the peer is gone down. This notification will cause the backup ILMs associated with this PG-ID to open and accept mLDP packets immediately. When IGP re-converges, an updated pair of primary and backup ILMs is downloaded for each mLDP FEC by LDP into the IOM with the corresponding PG-IDs.

If multiple LDP interfaces exist to the upstream LSR, a failure of one interface will bring down the link Hello adjacency on that interface but not the LDP session which is still associated with the remaining link Hello adjacencies. In this case, the upstream LSR updates in IOM the NHLFE for the mLDP FEC to use one of the remaining links. The switchover time in this case is not managed by the uniform failover procedures.

## 5.19 Multi-Area and Multi-Instance Extensions to LDP

In order to extend LDP across multiple areas of an IGP instance or across multiple IGP instances, the current standard LDP implementation based on RFC 3036 requires that all /32 prefixes of PEs be leaked between the areas or instances. This is because an exact match of the prefix in the routing table is required to install the prefix binding in the LDP Forwarding Information Base (FIB). Although a router will do this by default when configured as Area Border Router (ABR), this increases the convergence of IGP on routers when the number of PE nodes scales to thousands of nodes.

Multi-area and multi-instance extensions to LDP provide an optional behavior by which LDP installs a prefix binding in the LDP FIB by simply performing a longest prefix match with an aggregate prefix in the routing table (RIB). That way, the ABR will be configured to summarize the /32 prefixes of PE routers. This method is compliant to RFC 5283, *LDP Extension for Inter-Area Label Switched Paths (LSPs)*.

## 5.19.1 LDP Shortcut for BGP Next-Hop Resolution

LDP shortcut for BGP next-hop resolution shortcuts allow for the deployment of a 'route-less core' infrastructure on the 7750 SR and 7950 XRS. Many service providers either have or intend to remove the IBGP mesh from their network core, retaining only the mesh between routers connected to areas of the network that require routing to external routes.

Shortcuts are implemented by utilizing Layer 2 tunnels (i.e., MPLS LSPs) as next hops for prefixes that are associated with the far end termination of the tunnel. By tunneling through the network core, the core routers forwarding the tunnel have no need to obtain external routing information and are immune to attack from external sources.

The tunnel table contains all available tunnels indexed by remote destination IP address. LSPs derived from received LDP /32 route FECs will automatically be installed in the table associated with the advertising router-ID when IGP shortcuts are enabled.

Evaluating tunnel preference is based on the following order in descending priority:

1. LDP /32 route FEC shortcut
2. Actual IGP next-hop

If a higher priority shortcut is not available or is not configured, a lower priority shortcut is evaluated. When no shortcuts are configured or available, the IGP next-hop is always used. Shortcut and next-hop determination is event driven based on dynamic changes in the tunneling mechanisms and routing states.

Refer to the OS Routing Protocols Guide for details on the use of LDP FEC and RSVP LSP for BGP Next-Hop Resolution.

## 5.19.2 LDP Shortcut for IGP Routes

The LDP shortcut for IGP route resolution feature allows forwarding of packets to IGP learned routes using an LDP LSP. When LDP shortcut is enabled globally, IP packets forwarded over a network IP interface will be labeled with the label received from the next-hop for the route and corresponding to the FEC-prefix matching the destination address of the IP packet. In such a case, the routing table will have the shortcut next-hop as the best route. If such a LDP FEC does not exist, then the routing table will have the regular IP next-hop and regular IP forwarding will be performed on the packet.

An egress LER advertises and maintains a FEC, label binding for each IGP learned route. This is performed by the existing LDP fec-originate capability.

### 5.19.2.1 LDP Shortcut Configuration

The user enables the use of LDP shortcut for resolving IGP routes by entering the global command **config>router>ldp-shortcut**.

This command enables forwarding of user IP packets and specified control IP packets using LDP shortcuts over all network interfaces in the system which participate in the IS-IS and OSPF routing protocols. The default is to disable the LDP shortcut across all interfaces in the system.

### 5.19.2.2 IGP Route Resolution

When LDP shortcut is enabled, LDP populates the RTM with next-hop entries corresponding to all prefixes for which it activated an LDP FEC. For a given prefix, two route entries are populated in RTM. One corresponds to the LDP shortcut next-hop and has an owner of LDP. The other one is the regular IP next-hop. The LDP shortcut next-hop always has preference over the regular IP next-hop for forwarding user packets and specified control packets over a given outgoing interface to the route next-hop.

The prior activation of the FEC by LDP is done by performing an exact match with an IGP route prefix in RTM. It can also be done by performing a longest prefix-match with an IGP route in RTM if the aggregate-prefix-match option is enabled globally in LDP.

This feature is not restricted to /32 FEC prefixes. However only /32 FEC prefixes will be populated in the CPM Tunnel Table for use as a tunnel by services.

All user packets and specified control packets for which the longest prefix match in RTM yields the FEC prefix will be forwarded over the LDP LSP. Currently, the control packets that could be forwarded over the LDP LSP are ICMP ping and UDP-traceroute. The following is an example of the resolution process.

Assume the egress LER advertised a FEC for some /24 prefix using the fec-originate command. At the ingress LER, LDP resolves the FEC by checking in RTM that an exact match exists for this prefix. Once LDP activated the FEC, it programs the NHLFE in the egress data path and the LDP tunnel information in the ingress data path tunnel table.

Next, LDP provides the shortcut route to RTM which will associate it with the same /24 prefix. There will be two entries for this /24 prefix, the LDP shortcut next-hop and the regular IP next-hop. The latter was used by LDP to validate and activate the FEC. RTM then resolves all user prefixes which succeed a longest prefix match against the /24 route entry to use the LDP LSP.

Assume now the aggregate-prefix-match was enabled and that LDP found a /16 prefix in RTM to activate the FEC for the /24 FEC prefix. In this case, RTM adds a new more specific route entry of /24 and has the next-hop as the LDP LSP but it will still not have a specific /24 IP route entry. RTM then resolves all user prefixes which succeed a longest prefix match against the /24 route entry to use the LDP LSP while all other prefixes which succeed a longest prefix-match against the /16 route entry will use the IP next-hop.

### 5.19.2.3 LDP Shortcut Forwarding Plane

Once LDP activated a FEC for a given prefix and programmed RTM, it also programs the ingress Tunnel Table in forwarding engine with the LDP tunnel information.

When an IPv4 packet is received on an ingress network interface, or a subscriber IES interface, or a regular IES interface, the lookup of the packet by the ingress forwarding engine will result in the packet being sent labeled with the label stack corresponding to the NHLFE of the LDP LSP when the preferred RTM entry corresponds to an LDP shortcut.

If the preferred RTM entry corresponds to an IP next-hop, the IPv4 packet is forwarded unlabeled.

## 5.19.3 ECMP Considerations

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

When the preferred RTM entry corresponds to an LDP shortcut route, spraying will be performed across the multiple next-hops for the LDP FEC. The FEC next-hops can either be direct link LDP neighbors or T-LDP neighbors reachable over RSVP LSPs in the case of LDP-over-RSVP but not both. This is as per ECMP for LDP in existing implementation.

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

## 5.19.4 Disabling TTL Propagation in an LSP Shortcut

This feature provides the option for disabling TTL propagation from a transit or a locally generated IP packet header into the LSP label stack when an LDP LSP is used as a shortcut for BGP next-hop resolution, a static-route next-hop resolution, or for an IGP route resolution.

A transit packet is a packet received from an IP interface and forwarded over the LSP shortcut at ingress LER.

A locally-generated IP packet is any control plane packet generated from the CPM and forwarded over the LSP shortcut at ingress LER.

TTL handling can be configured for all LDP LSP shortcuts originating on an ingress LER using the following global commands:

```
config>router>ldp>[no] shortcut-transit-ttl-propagate
```

```
config>router>ldp>[no] shortcut-local-ttl-propagate
```

These commands apply to all LDP LSPs which are used to resolve static routes, BGP routes, and IGP routes.

When the **no** form of the above command is enabled for local packets, TTL propagation is disabled on all locally generated IP packets, including ICMP Ping, traceroute, and OAM packets that are destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack. This is referred to as pipe mode.

Similarly, when the **no** form is enabled for transit packets, TTL propagation is disabled on all IP packets received on any IES interface and destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack.

## 5.20 LDP Graceful Handling of Resource Exhaustion

This feature enhances the behavior of LDP when a data path or a CPM resource required for the resolution of a FEC is exhausted. In prior releases, the LDP module shuts down. The user is required to fix the issue causing the FEC scaling to be exceeded and to restart the LDP module by executing the **unshut** command.

## 5.20.1 LDP Base Graceful Handling of Resources

This feature implements a base graceful handling capability by which the LDP interface to the peer, or the targeted peer in the case of Targeted LDP (T-LDP) session, is shutdown. If LDP tries to resolve a FEC over a link or a targeted LDP session and it runs out of data path or CPM resources, it will bring down that interface or targeted peer which will bring down the Hello adjacency over that interface to the resolved link LDP peer or to the targeted peer. The interface is brought down in LDP context only and is still available to other applications such as IP forwarding and RSVP LSP forwarding.

Depending of what type of resource was exhausted, the scope of the action taken by LDP will be different. Some resource such as NHLFE have interface local impact, meaning that only the interface to the downstream LSR which advertised the label is shutdown. Some resources such as ILM have global impact, meaning that they will impact every downstream peer or targeted peer which advertised the FEC to the node. The following are examples to illustrate this.

- For NHLFE exhaustion, one or more interfaces or targeted peers, if the FEC is ECMP, will be shut down. ILM is maintained as long as there is at least one downstream for the FEC for which the NHLFE has been successfully programmed.
- For an exhaustion of an ILM for a unicast LDP FEC, all interfaces to peers or all target peers which sent the FEC will be shutdown. No deprogramming of data path is required since FEC is not programmed.
- An exhaustion of ILM for an mLDP FEC can happen during primary ILM programming, MBB ILM programming, or multicast upstream FRR backup ILM programming. In all cases, the P2MP index for the mLDP tree is deprogrammed and the interfaces to each downstream peer which sent a Label Mapping message associated with this ILM are shutdown.

After the user has taken action to free resources up, he/she will require manually unshut the interface or the targeted peer to bring it back into operation. This then re-establishes the Hello adjacency and resumes the resolution of FECs over the interface or to the targeted peer.

Detailed guidelines for using the feature and for troubleshooting a system which activated this feature are provided in the following sections.

This new behavior will become the new default behavior in Release 11.0.R4 and will interoperate with SR OS based LDP implementation and any other third party LDP implementation.

The following data path resources can trigger this mechanism:



- NHLFE
- ILM
- Label-to-NHLFE (LTN)
- Tunnel Index
- P2MP Index

The following CPM resources can trigger this mechanism:

- Label allocation

## 5.21 LDP Enhanced Graceful Handling of Resources

This feature is an enhanced graceful handling capability which is supported only among SR OS based implementations. If LDP tries to resolve a FEC over a link or a targeted session and it runs out of data path or CPM resources, it will put the LDP/T-LDP session into overload state. As a result, it will release to its LDP peer the labels of the FECs which it could not resolve and will also send an LDP notification message to all LDP peers with the new status load of overload for the FEC type which caused the overload. The notification of overload is per FEC type, i.e., unicast IPv4, P2MP mLDP etc., and not per individual FEC. The peer which caused the overload and all other peers will stop sending any new FECs of that type until this node updates the notification stating that it is no longer in overload state for that FEC type. FECs of this type previously resolved and other FEC types to this peer and all other peers will continue to forward traffic normally.

After the user has taken action to free resources up, he/she will require manually clear the overload state of the LDP/T-LDP sessions towards its peers.

The enhanced mechanism will be enabled instead of the base mechanism only if both LSR nodes advertise this new LDP capability at the time the LDP session is initialized. Otherwise, they will continue to use the base mechanism.

This feature operates among SR OS LSR nodes using a couple of private vendor LDP capabilities:

- The first one is the LSR Overload Status TLV to signal or clear the overload condition.
- The second one is the Overload Protection Capability Parameter which allows LDP peers to negotiate the use or not of the overload notification feature and hence the enhanced graceful handling mechanism.

When interoperating with an LDP peer which does not support the enhanced resource handling mechanism, the router reverts automatically to the default base resource handling mechanism.

The following are the details of the mechanism.

### 5.21.1 LSR Overload Notification

When an upstream LSR is overloaded for a FEC type, it notifies one or more downstream peer LSRs that it is overloaded for the FEC type.

When a downstream LSR receives overload status ON notification from an upstream LSR, it does not send further label mappings for the specified FEC type. When a downstream LSR receives overload OFF notification from an upstream LSR, it sends pending label mappings to the upstream LSR for the specified FEC type.

This feature introduces a new TLV referred to as *LSR Overload Status TLV*, shown below. This TLV is encoded using vendor proprietary TLV encoding as per RFC 5036. It uses a TLV type value of 0x3E02 and the Timetra OUI value of 0003FA.

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|U|F| Overload Status TLV Type | Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Timetra OUI = 0003FA |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|S| Reserved |

```

where:

U-bit: Unknown TLV bit, as described in RFC 5036. The value MUST be 1 which means if unknown to receiver then receiver should ignore

F-bit: Forward unknown TLV bit, as described in RFC RFC5036. The value of this bit MUST be 1 since a LSR overload TLV is sent only between two immediate LDP peers, which are not forwarded.

S-bit: The State Bit. It indicates whether the sender is setting the LSR Overload Status ON or OFF. The State Bit value is used as follows:

1 - The TLV is indicating LSR overload status as ON.

0 - The TLV is indicating LSR overload status as OFF.

When a LSR that implements the procedures defined in this document generates LSR overload status, it **MUST** send LSR Overload Status TLV in a LDP Notification Message accompanied by a FEC TLV. The FEC TLV must contain one Typed Wildcard FEC TLV that specifies the FEC type to which the overload status notification applies.

The feature in this document re-uses the Typed Wildcard FEC Element which is defined in RFC 5918.

## 5.21.2 LSR Overload Protection Capability

To ensure backward compatibility with procedures in RFC 5036 an LSR supporting Overload Protection need means to determine whether a peering LSR supports overload protection or not.

An LDP speaker that supports the LSR Overload Protection procedures as defined in this document **MUST** inform its peers of the support by including a LSR Overload Protection Capability Parameter in its initialization message. The Capability parameter follows the guidelines and all Capability Negotiation Procedures as defined in RFC 5561. This TLV is encoded using vendor proprietary TLV encoding as per RFC 5036. It uses a TLV type value of 0x3E03 and the Timetra OUI value of 0003FA.

```

 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
 +-+
 |U|F| LSR Overload Cap TLV Type | |
 +-+
 | Timetra OUI = 0003FA |
 +-+
 |S| Reserved |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

```

Where:

U and F bits : **MUST** be 1 and 0 respectively as per section 3 of LDP Capabilities [RFC5561].

S-bit : **MUST** be 1 (indicates that capability is being advertised).

## 5.21.3 Procedures for LSR overload protection

The procedures defined in this document apply only to LSRs that support Downstream Unsolicited (DU) label advertisement mode and Liberal Label Retention Mode. An LSR that implements the LSR overload protection follows the following procedures:

1. A LSR MUST NOT use LSR Overload notification procedures with a peer LSR that has not specified LSR Overload Protection Capability in Initialization Message received from the peer LSR.
2. When an upstream LSR detects that it is overloaded with a FEC type then it MUST initiate a LDP Notification Message with S bit ON in LSR Overload Status TLV and a FEC TLV containing the Typed Wildcard FEC Element for the specified FEC type. The Message may be sent to one or more peers.
3. After it has notified overload status ON for a FEC type, the overloaded upstream LSR MAY send Label Release for a set of FEC elements to respective downstream LSRs to off load its LIB below certain watermark.
4. When an upstream LSR that was overloaded for a FEC type before, detects that it is no longer overloaded then it MUST send a LDP Notification Message with S bit OFF in LSR Overload Status TLV and FEC TLV containing the Typed Wildcard FEC Element for the specified FEC type.
5. When an upstream LSR has notified as overloaded for a FEC type, then a downstream LSR MUST NOT send new Label Mappings for the specified FEC type to the upstream LSR.
6. When a downstream LSR receives LSR Overload Notification from a peering LSR with status OFF for a FEC type then the receiving LSR MUST send any label mappings for the FEC type which were pending to the upstream LSR or which are eligible to be sent now.
7. When an upstream LSR is overloaded for a FEC type and it receives Label Mapping for that FEC type from a downstream LSR then it MAY send Label Release to the downstream for the received Label Mapping with LDP Status Code as *No\_Label\_Resources* as defined in RFC 5036.

## 5.22 User Guidelines and Troubleshooting Procedures

### 5.22.1 Common Procedures

When troubleshooting a LDP resource exhaustion situation on an LSR, the user must first determine which of the LSR and its peers supports the enhanced handling of resources. This is done by checking if the local LSR or its peers advertised the LSR Overload Protection Capability:

```
show router ldp status
=====
LDP Status for LSR ID 110.20.1.110
```

```

=====
Admin State : Up Oper State : Up
Created at : 07/17/13 21:27:41 Up Time : 0d 01:00:41
Oper Down Reason : n/a Oper Down Events : 1
Last Change : 07/17/13 21:27:41 Tunn Down Damp Time : 20 sec
Label Withdraw Del*: 0 sec Implicit Null Label : Enabled
Short. TTL Prop Lo*: Enabled Short. TTL Prop Tran*: Enabled
Import Policies : Export Policies :
 Import-LDP Import-LDP
 External External
Tunl Exp Policies :
 from-proto-bgp
Aggregate Prefix : False Agg Prefix Policies : None
FRR : Enabled Mcast Upstream FRR : Disabled
Dynamic Capability : False P2MP Capability : True
MP MBB Capability : True MP MBB Time : 10
Overload Capability: True <---- //Local Overload Capability
Active Adjacencies : 0 Active Sessions : 0
Active Interfaces : 2 Inactive Interfaces : 4
Active Peers : 62 Inactive Peers : 10
Addr FECs Sent : 0 Addr FECs Recv : 0
Serv FECs Sent : 0 Serv FECs Recv : 0
P2MP FECs Sent : 0 P2MP FECs Recv : 0
Attempted Sessions : 458
No Hello Err : 0 Param Adv Err : 0
Max PDU Err : 0 Label Range Err : 0
Bad LDP Id Err : 0 Bad PDU Len Err : 0
Bad Mesg Len Err : 0 Bad TLV Len Err : 0
Unknown TLV Err : 0
Malformed TLV Err : 0 Keepalive Expired Err: 4
Shutdown Notif Sent: 12 Shutdown Notif Recv : 5
=====

```

show router ldp session detail

LDP Sessions (Detail)

-----  
Session with Peer 10.8.100.15:0, Local 110.20.1.110:0  
-----

```

Adjacency Type : Targeted State : Nonexistent
Up Time : 0d 00:00:00
Max PDU Length : 4096 KA/Hold Time Remaining : 0
Link Adjacencies : 0 Targeted Adjacencies : 1
Local Address : 110.20.1.110 Peer Address : 10.8.100.15
Local TCP Port : 0 Peer TCP Port : 0
Local KA Timeout : 40 Peer KA Timeout : 40
Mesg Sent : 0 Mesg Recv : 1
FECs Sent : 0 FECs Recv : 0
Addrs Sent : 0 Addrs Recv : 0
GR State : Capable Label Distribution : DU
Nbr Liveness Time : 0 Max Recovery Time : 0
Number of Restart : 0 Last Restart Time : Never
P2MP : Not Capable MP MBB : Not Capable
Dynamic Capability : Not Capable LSR Overload : Not Capable <----
//Peer OverLoad Capab.
Advertise : Address/Servi*
Addr FEC OverLoad Sent : No Addr FEC OverLoad Recv : No
Mcast FEC Overload Sent: No Mcast FEC Overload Recv: No

```

Serv FEC Overload Sent : No

Serv FEC Overload Recv : No

## 5.22.2 Base Resource Handling Procedures

### Step 1

If the peer OR the local LSR does not support the Overload Protection Capability it means that the associated adjacency [interface/peer] will be brought down as part of the base resource handling mechanism.

The user can determine which interface or targeted peer was shut down, by applying the following commands:

- [show router ldp interface resource-failures]

- [show router ldp targ-peer resource-failures]

```
show router ldp interface resource-failures
=====
LDP Interface Resource Failures
=====
srl srr
sru4 sr4-1-5-1
=====

show router ldp targ-peer resource-failures
=====
LDP Peers Resource Failures
=====
10.20.1.22 110.20.1.3
=====
```

A trap is also generated for each interface or targeted peer:

```
16 2013/07/17 14:21:38.06 PST MINOR: LDP #2003 Base LDP Interface Admin State
"Interface instance state changed - vRtrID: 1, Interface sr4-1-5-1, administrative state: inService, operational state: outOfService"

13 2013/07/17 14:15:24.64 PST MINOR: LDP #2003 Base LDP Interface Admin State
"Interface instance state changed - vRtrID: 1, Peer 10.20.1.22, administrative state: inService, operational state: outOfService"
```

The user can then check that the base resource handling mechanism has been applied to a specific interface or peer by running the following show commands:

- [show router ldp interface detail]

- [show router ldp targ-peer detail]

```

show router ldp interface detail
=====
LDP Interfaces (Detail)
=====

Interface "sr4-1-5-1"

Admin State : Up Oper State : Down
Oper Down Reason : noResources <----- //link LDP resource exhaustion handled
Hold Time : 45 Hello Factor : 3
Oper Hold Time : 45
Hello Reduction : Disabled Hello Reduction *: 3
Keepalive Timeout : 30 Keepalive Factor : 3
Transport Addr : System Last Modified : 07/17/13 14:21:38
Active Adjacencies : 0
Tunneling : Disabled
Lsp Name : None
Local LSR Type : System
Local LSR : None
BFDD Status : Disabled
Multicast Traffic : Enabled

show router ldp discovery interface "sr4-1-5-1" detail
=====
LDP Hello Adjacencies (Detail)
=====

Interface "sr4-1-5-1"

Local Address : 223.0.2.110 Peer Address : 224.0.0.2
Adjacency Type : Link State : Down
=====

show router ldp targ-peer detail
=====
LDP Peers (Detail)
=====

Peer 10.20.1.22

Admin State : Up Oper State : Down
Oper Down Reason : noResources <----- // T-LDP resource exhaustion handled
Hold Time : 45 Hello Factor : 3
Oper Hold Time : 45
Hello Reduction : Disabled Hello Reduction Fact*: 3
Keepalive Timeout : 40 Keepalive Factor : 4
Passive Mode : Disabled Last Modified : 07/17/13 14:15:24
Active Adjacencies : 0 Auto Created : No
Tunneling : Enabled
Lsp Name : None
Local LSR : None
BFDD Status : Disabled
Multicast Traffic : Disabled

show router ldp discovery peer 10.20.1.22 detail

```

```

=====
LDP Hello Adjacencies (Detail)
=====

Peer 10.20.1.22

Local Address : 110.20.1.110 Peer Address : 10.20.1.22
Adjacency Type : Targeted State : Down <-----
//T-LDP resource exhaustion handled
=====

```

## Step 2

Besides interfaces and targeted peer, locally originated FECs may also be put into overload. These are the following:

- unicast fec-originate pop
- multicast local static p2mp-fec type=1 [on leaf LSR]
- multicast local Dynamic p2mp-fec type=3 [on leaf LSR]

The user can check if only remote and/or local FECs have been set in overload by the resource base resource exhaustion mechanism using the following command:

- [tools dump router ldp instance]

The relevant part of the output is described below:

```

{..... snip.....}
Num OLoad Interfaces: 4 <----- //LDP interfaces resource in exhaustion
Num Targ Sessions: 72 Num Active Targ Sess: 62
Num OLoad Targ Sessions: 7 <----- //T-LDP peers in resource exhaustion
Num Addr FECs Rcvd: 0 Num Addr FECs Sent: 0
Num Addr Fecs OLoad: 1 <----- // # of local/remote unicast FECs in Overload
Num Svc FECs Rcvd: 0 Num Svc FECs Sent: 0
Num Svc FECs OLoad: 0 <----- // # of local/
remote service Fecs in Overload
Num mcast FECs Rcvd: 0 Num Mcast FECs Sent: 0
Num mcast FECs OLoad: 0 <----- // # of local/
remote multicast Fecs in Overload
{..... snip.....}

```

When at least one local FEC has been set in overload the following trap will occur:

```

23 2013/07/
17 15:35:47.84 PST MINOR: LDP #2002 Base LDP Resources Exhausted "Instance
state changed - vRtrID: 1, administrative state: inService, operationa l state:
inService"

```

## Step 3



After the user has detected that at least, one link LDP or T-LDP adjacency has been brought down by the resource exhaustion mechanism, he/she must protect the router by applying one or more of the following to free resources up:

- Identify the source for the [unicast/multicast/service] FEC flooding.
- Configure the appropriate [import/export] policies and/or delete the excess [unicast/multicast/service] FECs not currently handled.

#### Step 4

Next, the user has to manually attempt to clear the overload (no resource) state and allow the router to attempt to restore the link and targeted sessions to its peer.



**Note:** Because of the dynamic nature of FEC distribution and resolution by LSR nodes, one cannot predict exactly which FECs and which interfaces or targeted peers will be restored after performing the following commands if the LSR activates resource exhaustion again.

One of the following commands can be used:

- [**clear router ldp resource-failures**]

- Clears the overload state and attempt to restore adjacency and session for LDP interfaces and peers.
- Clear the overload state for the local FECs.

- [**clear router ldp interface ifName**]

- [**clear router ldp peer peerAddress**]

- Clears the overload state and attempt to restore adjacency and session for LDP interfaces and peers.
- These two commands **\*DO NOT\*** Clear the overload state for the local FECs.

## 5.22.3 Enhanced Resource Handling Procedures

#### Step 1

If the peer AND the local LSR do support the Overload Protection Capability it means that the LSR will signal the overload state for the FEC type which caused the resource exhaustion as part of the enhanced resource handling mechanism.

In order to verify if the local router has received or sent the overload status TLV, perform the following:

```

-[show router ldp session detail]
show router ldp session 110.20.1.1 detail

Session with Peer 110.20.1.1:0, Local 110.20.1.110:0

Adjacency Type : Both State : Established
Up Time : 0d 00:05:48
Max PDU Length : 4096 KA/Hold Time Remaining : 24
Link Adjacencies : 1 Targeted Adjacencies : 1
Local Address : 110.20.1.110 Peer Address : 110.20.1.1
Local TCP Port : 51063 Peer TCP Port : 646
Local KA Timeout : 30 Peer KA Timeout : 45
Mesg Sent : 442 Mesg Recv : 2984
FECs Sent : 16 FECs Recv : 2559
Addrs Sent : 17 Addrs Recv : 1054
GR State : Capable Label Distribution : DU
Nbr Liveness Time : 0 Max Recovery Time : 0
Number of Restart : 0 Last Restart Time : Never
P2MP : Capable MP MBB : Capable
Dynamic Capability : Not Capable LSR Overload : Capable
Advertise : Address/Servi* BFD Operational Status : inService
Addr FEC OverLoad Sent : Yes Addr FEC OverLoad Recv : No <----
// this LSR sent overLoad for unicast FEC type to peer
Mcast FEC Overload Sent: No Mcast FEC Overload Recv: No
Serv FEC Overload Sent : No Serv FEC Overload Recv : No

show router ldp session 110.20.1.110 detail

Session with Peer 110.20.1.110:0, Local 110.20.1.1:0

Adjacency Type : Both State : Established
Up Time : 0d 00:08:23
Max PDU Length : 4096 KA/Hold Time Remaining : 21
Link Adjacencies : 1 Targeted Adjacencies : 1
Local Address : 110.20.1.1 Peer Address : 110.20.1.110
Local TCP Port : 646 Peer TCP Port : 51063
Local KA Timeout : 45 Peer KA Timeout : 30
Mesg Sent : 3020 Mesg Recv : 480
FECs Sent : 2867 FECs Recv : 16
Addrs Sent : 1054 Addrs Recv : 17
GR State : Capable Label Distribution : DU
Nbr Liveness Time : 0 Max Recovery Time : 0
Number of Restart : 0 Last Restart Time : Never
P2MP : Capable MP MBB : Capable
Dynamic Capability : Not Capable LSR Overload : Capable
Advertise : Address/Servi* BFD Operational Status : inService
Addr FEC OverLoad Sent : No Addr FEC OverLoad Recv : Yes <----
// this LSR received overLoad for unicast FEC type from peer
Mcast FEC Overload Sent: No Mcast FEC Overload Recv: No
Serv FEC Overload Sent : No Serv FEC Overload Recv : No
=====

```

A trap is also generated:

```

70002 2013/07/17 16:06:59.46 PST MINOR: LDP #2008 Base LDP Session State Change
"Session state is operational. Overload Notification message is sent to/from peer
 110.20.1.1:0 with overload state true for fec type prefixes"

```

## Step 2

Besides interfaces and targeted peer, locally originated FECs may also be put into overload. These are the following:

- unicast fec-originate pop
- multicast local static p2mp-fec type=1 [on leaf LSR]
- multicast local Dynamic p2mp-fec type=3 [on leaf LSR]

The user can check if only remote and/or local FECs have been set in overload by the resource enhanced resource exhaustion mechanism using the following command:

- **[tools dump router ldp instance]**

The relevant part of the output is described below:

```

 Num Entities OLoad (FEC: Address Prefix): Sent: 7 Rcvd: 0 <-----
// # of session in OvLd for fec-type=unicast
 Num Entities OLoad (FEC: PWE3): Sent: 0 Rcvd: 0 <-----
// # of session in OvLd for fec-type=service
 Num Entities OLoad (FEC: GENPWE3): Sent: 0 Rcvd: 0 <-----
// # of session in OvLd for fec-type=service
 Num Entities OLoad (FEC: P2MP): Sent: 0 Rcvd: 0 <-----
// # of session in OvLd for fec-type=MulticastP2mp
 Num Entities OLoad (FEC: MP2MP UP): Sent: 0 Rcvd: 0 <-----
// # of session in OvLd for fec-type=MulticastMP2mp
 Num Entities OLoad (FEC: MP2MP DOWN): Sent: 0 Rcvd: 0 <-----
// # of session in OvLd for fec-type=MulticastMP2mp
 Num Active Adjacencies: 9
 Num Interfaces: 6 Num Active Interfaces: 6
 Num OLoad Interfaces: 0 <----- //
link LDP interfaces in resource exhaustion
should be zero when Overload Protection Capability is supported
 Num Targ Sessions: 72 Num Active Targ Sess: 67
 Num OLoad Targ Sessions: 0 <----- // T-LDP peers in resource exhaustion
should be zero if Overload Protection Capability is supported
 Num Addr FECs Rcvd: 8667 Num Addr FECs Sent: 91
 Num Addr Fecs OLoad: 1 <-----
// # of local/remote unicast Fecs in Overload
 Num Svc FECs Rcvd: 3111 Num Svc FECs Sent: 0
 Num Svc FECs OLoad: 0 <-----
// # of local/remote service Fecs in Overload
 Num mcast FECs Rcvd: 0 Num Mcast FECs Sent: 0
 Num mcast FECs OLoad: 0 <-----
// # of local/remote multicast Fecs in Overload
 Num MAC Flush Rcvd: 0 Num MAC Flush Sent: 0

```

When at least one local FEC has been set in overload the following trap will occur:

```

69999 2013/07/17 16:06:59.21 PST MINOR: LDP #2002 Base LDP Resources Exhausted
"Instance state changed - vRtrID: 1, administrative state: inService, operational

```

```
state: inService"
```

### Step 3

After the user has detected that at least one overload status TLV has been sent or received by the LSR, he/she must protect the router by applying one or more of the following to free resources up:

- Identify the source for the [unicast/multicast/service] FEC flooding. This is most likely the LSRs which session received the overload status TLV.
- Configure the appropriate [import/export] policies and/or delete the excess [unicast/multicast/service] FECs from the FEC type in overload.

### Step 4

Next, the user has to manually attempt to clear the overload state on the affected sessions and for the affected FEC types and allow the router to clear the overload status TLV to its peers.



**Note:** Because of the dynamic nature of FEC distribution and resolution by LSR nodes, one cannot predict exactly which sessions and which FECs will be cleared after performing the following commands if the LSR activates overload again.

One of the following commands can be used depending if the user wants to clear all sessions or at once or one session at a time:

#### - [clear router ldp resource-failures]

- Clears the overload state for the affected sessions and FEC types.
- Clear the overload state for the local FECs.

#### - [clear router ldp session a.b.c.d overload fec-type {services | prefixes | multicast}]

- Clears the overload state for the specified session and FEC type.
- Clears the overload state for the local FECs.

## 5.23 LDP IPv6 Control and Data Planes


SR OS extends the LDP control plane and data plane to support LDP IPv6 adjacency and session using 128-bit LSR-ID.

The implementation allows for concurrent support of independent LDP IPv4 (32-bit LSR-ID) and IPv6 (128-bit LSR-ID) adjacencies and sessions between peer LSRs and over the same or different set of interfaces.

## 5.23.1 LDP Operation in an IPv6 Network

LDP IPv6 can be enabled on the SR OS interface. [Figure 75](#) shows the LDP adjacency and session over an IPv6 interface.

**Figure 75 LDP Adjacency and Session over an IPv6 Interface**



LSR-A and LSR-B have the following IPv6 LDP identifiers respectively:

- <LSR Id=A/128> : <label space id=0>
- <LSR Id=B/128> : <label space id=0>

By default, A/128 and B/128 use the system interface IPv6 address.



**Note:** Although the LDP control plane can operate using only the IPv6 system address, the user must configure the IPv4-formatted router ID for OPSF, IS-IS, and BGP to operate properly.

The following sections describe the behavior when LDP IPv6 is enabled on the interface.

## 5.23.2 Link LDP

The SR OS LDP IPv6 implementation uses a 128-bit LSR-ID as defined in *draft-pdutta-mpls-ldp-v2-00*. See [LDP Process Overview](#) for more information about interoperability of this implementation with 32-bit LSR-ID, as defined in *draft-ietf-mpls-ldp-ipv6-14*.

Hello adjacency will be brought up using link Hello packet with 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, will be set to the LSR-ID of the LSR by default. It will be set to the interface IPv6 address if the user enabled the interface option under one of the following contexts:

- **config>router>ldp>if-params>ipv6>transport-address**
- **config>router>ldp>if-params>if>ipv6>transport-address**

The interface global unicast address, meaning the primary IPv6 unicast address of the interface, is used.

The user can configure the **local-lsr-id** option on the interface and change the value of the LSR-ID to either the local interface or to another interface name, 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 which 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 will not come up and an error message will be displayed.

The LSR with the highest transport address will bootstrap the IPv6 TCP connection and IPv6 LDP session.

Source and destination addresses of LDP/TCP session packets are the IPv6 transport addresses.

### 5.23.3 Targeted LDP

Source and destination addresses of targeted Hello packet are the LDP IPv6 LSR-IDs of systems A and B.

The user can configure the **local-lsr-id** option on the targeted session and change the value of the LSR-ID to either the local interface or to some other interface name, loopback or not. 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 which 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 will not come up and an error message will be displayed. In all cases, the transport address for the LDP session and the source IP address of targeted Hello message will be updated to the new LSR-ID value.

The LSR with the highest transport address (in this case, the LSR-ID) will bootstrap the IPv6 TCP connection and IPv6 LDP session.

Source and destination IP addresses of LDP/TCP session packets are the IPv6 transport addresses (in this case, LDP LSR-IDs of systems A and B).

## 5.23.4 FEC Resolution

LDP will advertise and withdraw 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.

All LDP FEC types can be exchanged over a LDP IPv6 LDP session like in LDP IPv4 session.

The LSR does not advertise a FEC for a link-local address and, if received, the LSR will not resolve it.

A IPv4 or IPv6 prefix FEC can be resolved to an LDP IPv6 interface in the same way as it is resolved to an LDP IPv4 interface. The outgoing interface and next-hop are looked up in RTM cache. The next-hop can be the link-local unicast address of the other side of the link or a global unicast address. The FEC is resolved to the LDP IPv6 interface of the downstream LDP IPv6 LSR that advertised the IPv4 or IPv6 address of the next hop.

An mLDP P2MP FEC with an IPv4 root LSR address, and carrying one or more IPv4 or IPv6 multicast prefixes in the opaque element, can be resolved to an upstream LDP IPv6 LSR by checking if the LSR advertised the next-hop for the IPv4 root LSR address. The upstream LDP IPv6 LSR will then resolve the IPv4 P2MP FEC to one of the LDP IPV6 links to this LSR.



**Note:** Beginning in Release 13.0, a P2MP FEC with an IPv6 root LSR address, carrying one or more IPv4 or IPv6 multicast prefixes in the opaque element, is not supported. Manually configured mLDP P2MP LSP, NG-mVPN, and dynamic mLDP will not be able to operate in an IPv6-only network.

A PW FEC can be resolved to a targeted LDP IPv6 adjacency with an LDP IPv6 LSR if there is a context for the FEC with local spoke-SDP configuration or spoke-SDP auto-creation from a service such as BGP-AD VPLS, BGP-VPWS or dynamic MS-PW.

## 5.23.5 LDP Session Capabilities

LDP supports advertisement of all FEC types over an LDP IPv4 or an LDP IPv6 session. These FEC types are: IPv4 prefix FEC, IPv6 prefix FEC, IPv4 P2MP FEC, PW FEC 128, and PW FEC 129.

In addition, LDP supports signaling the enabling or disabling of the advertisement of the following subset of FEC types both during the LDP IPv4 or IPv6 session initialization phase, and subsequently when the session is already up.

- IPv4 prefix FEC—This is performed using the State Advertisement Control (SAC) capability TLV as 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 a 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 a LDP Capability message to update the state of this FEC type.
- P2MP FEC—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. Unlike the IPv4 SAC and IPv6 SAC capabilities, the P2MP capability does not distinguish between IPv4 and IPv6 P2MP FEC. The LSR can send this TLV in the LDP Initialization message and, subsequently, in a LDP Capability message to update the state of this FEC type.

During LDP session initialization, each LSR indicates to its peers which FEC type it supports by including the capability TLV for it in the LDP Initialization message. The SR OS implementation will enable the above FEC types by default and will thus send the corresponding capability TLVs in the LDP initialization message. If one or both peers advertise the disabling of a capability in the LDP Initialization message, no FECs of the corresponding FEC type will be exchanged between the two peers for the lifetime of the LDP session unless a Capability message is sent subsequently to explicitly enable it. The same behavior applies if no capability TLV for a FEC type is advertised in the LDP initialization message, except for the IPv4 prefix FEC which is assumed to be supported by all implementations by default.



Dynamic Capability, as defined in RFC 5561, allows all above FEC types to update the enabled or disabled state after the LDP session initialization phase. An LSR informs its peer that it supports the Dynamic Capability by including the Dynamic Capability Announcement TLV in the LDP Initialization message. If both LSRs advertise this capability, the user is allowed to enable or disable any of the 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 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 will only take effect at the next time the LDP session is restarted.

The user can enable or disable a specific FEC type for a given LDP session to a peer by using the following CLI commands:

- **config>router>ldp>session-params>peer>fec-type-capability p2mp**
- **config>router>ldp>session-params>peer>fec-type-capability prefix-ipv4**
- **config>router>ldp>session-params>peer>fec-type-capability prefix-ipv6**

## 5.23.6 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 given FEC type is enabled at the session level, it can be disabled over a given LDP interface at the IPv4 or IPv6 adjacency level for all IPv4 or IPv6 peers over that interface. If a given FEC type is disabled at the session level, then FECs will not be advertised and enabling that FEC type at the adjacency level will 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 given 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 a LSR peer may be each be configured to carry one of the two FEC types. An mLDP P2MP FEC can exclude specific links to a downstream LSR from being used to resolve this type of FEC.

Like the LDP session-level FEC-type capability, the adjacency FEC-type capability is negotiated for both directions of the adjacency. If one or both peers advertise the disabling of a capability in the LDP Hello message, no FECs of the corresponding FEC type will be 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 specific FEC type for a given LDP interface to a peer by using the following CLI commands:

- **config>router>ldp>if-params>if>ipv4/ipv6>fec-type-capability p2mp-ipv4**
- **config>router>ldp>if-params>if>ipv4/ipv6>fec-type-capability p2mp-ipv6**
- **config>router>ldp>if-params>if>ipv4/ipv6>fec-type-capability prefix-ipv4**
- **config>router>ldp>if-params>if> ipv4/ipv6>fec-type-capability prefix-ipv6**

These commands, when applied for the P2MP FEC, deprecate the existing command **multicast-traffic {enable | disable}** under the interface. Unlike the session-level capability, these commands can disable multicast FEC for IPv4 and IPv6 separately.

The encoding of the adjacency capability TLV uses a PRIVATE Vendor TLV. It is used only in a hello message to negotiate a set of capabilities for a specific LDP IPv4 or IPv6 hello adjacency.

[illegible]

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; thus, 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

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.

## 5.23.7 Address and FEC Distribution

After an LDP LSR initializes the LDP session to the peer LSR and the session comes up, local IPv4 and IPv6 interface addresses are exchanged using the Address and Address Withdraw messages. Similarly, FECs are exchanged using Label Mapping messages.


By default, IPv6 address distribution is determined by whether the Dual-stack capability TLV, which is defined in *draft-ietf-mpls-ldp-ipv6*, is present in the Hello message from the peer. This coupling is introduced because of interoperability issues found with existing third-party LDP IPv4 implementations.

The following is the detailed behavior:


- If the peer sent the dual-stack capability TLV in the Hello message, then IPv6 local addresses will be sent to the peer. The user can configure a new address export policy to further restrict which local IPv6 interface addresses to send to the peer. If the peer explicitly stated enabling of LDP IPv6 FEC type by including the IPv6 SAC TLV with the D-bit (Disable-bit) set to 0 in the initialization message, then IPv6 FECs will be sent to the peer. FEC prefix export policies can be used to restrict which LDP IPv6 FEC can be 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 with the D-bit (Disable-bit) set to 1 in the initialization message, then IPv6 FECs will not be sent but IPv6 local addresses will be sent to the peer. A CLI is provided to allow the configuration of an address export policy to further restrict which local IPv6 interface addresses to send to the peer. FEC prefix export policy has no effect because the peer explicitly requested disabling the IPv6 FEC type advertisement.
- If the peer did not send the dual-stack capability TLV in the Hello message, then no IPv6 addresses or IPv6 FECs will be 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 existing third-party LDP IPv4 implementations. The user can override this by explicitly configuring an address export policy and a FEC export policy to select which 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 [Figure 76](#) and [Figure 77](#).

**Figure 76 LDP IPv6 Address and FEC Distribution Procedure**



al\_0625

**Figure 77 LDP IPv6 Address and FEC Distribution Procedure**

al\_0624

### 5.23.8 Controlling IPv6 FEC Distribution During an Upgrade to SR OS Supporting LDP IPv6

A FEC for each of the IPv4 and IPv6 system interface addresses is advertised and resolved automatically by the LDP peers when the LDP session comes up, regardless of whether the session is IPv4 or IPv6.

To avoid the automatic advertisement and resolution of IPv6 system FEC when the LDP session is IPv4, the following procedure must be followed before and after the upgrade to the SR OS version which introduces support of LDP IPv6.


1. Before the upgrade, implement a global prefix policy which rejects prefix [::0/0 longer] to prevent IPv6 FECs from being installed after the upgrade.
2. In MISSU case:
  - If new IPv4 sessions are created on the node, the per-peer FEC-capabilities must be configured to filter out IPv6 FECs.
  - Until an existing IPv4 session is flapped, FEC-capabilities have no effect on filtering out IPv6 FECs, thus the import global policy must remain configured in place until the session flaps. Alternatively, a per-peer-import-policy [::0/0 longer] can be associated with this peer.
3. In cold upgrade case:
  - If new IPv4 sessions are created on the node, the per-peer FEC-capabilities must be configured to filter out IPv6 FECs.
  - On older, pre-existing IPv4 sessions, the per-peer FEC-capabilities must be configured to filter out IPv6 FECs.
4. When all LDP IPv4 sessions have dynamic capabilities enabled, with per-peer FEC-capabilities for IPv6 FECs disabled, then the GLOBAL IMPORT policy can be removed.

### 5.23.9 Handling of Duplicate Link-Local IPv6 Addresses in FEC Resolution

Link-local IPv6 addresses are scoped to a link and, as such, duplicate addresses can be used on different links to the same or different peer LSR. When the duplicate addresses exist on the same LAN, routing will detect them and block one of them. In all other cases, duplicate links are valid because they are scoped to the local link.

In this section, LLn refers to Link-Local address (n).

[Figure 78](#) shows FEC resolution in a LAN.


**Figure 78 FEC Resolution in LAN**

LSR B resolves a mLDP FEC with the root node being Root LSR. The route lookup shows that best route to loopback of Root LSR is {interface if-B and next-hop LL1}.

However, LDP will find that both LSR A and LSR C advertised address LL1 and that there are hello adjacencies (IPv4 or IPv6) to both A and C. In this case, a change is made so that an LSR only advertises link-local IPv6 addresses to a peer for the links over which it established a Hello adjacency to that peer. In this case, LSR C will advertise LL1 to LSR E but not to LSRs A, B, and D. This behavior will apply with both P2P and broadcast interfaces.

Ambiguity also exists with prefix FEC (unicast FEC); the above solution also applies.

#### FEC Resolution over P2P links



LSR B resolves an mLDP FEC with root node being Root LSR. The route lookup shows that best route to loopback of Root LSR is {interface if-B and next-hop LL1}.

- Case 1—LDP is enabled on all links. This case has no ambiguity. LDP will only select LSR A because the address LL1 from LSR C is discovered over a different interface. This case also applies to prefix FEC (unicast FEC) and thus no ambiguity in the resolution.



- Case 2—LDP is disabled on link A-B with next-hop LL1; LSR B can still select one of the two other interfaces to upstream LSR A as long as LSR A advertised LL1 address in the LDP session.

## 5.23.10 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/IPv6 LDP interface as follows:

1. If the router interface goes down or both LDP IPv4 and LDP IPv6 sessions go down, IGP sets the interface metric to maximum value and all static routes with the **ldp-sync** option enabled and resolved on this interface will be de-activated.
2. If the router interface is up and only one of the LDP IPv4 or LDP IPv6 interfaces goes down, no action is taken.
3. When the router interface comes up from a down state, and one of either the LDP IPv4 or LDP IPv6 sessions comes up, IGP starts the sync timer at the expiry of which the interface metric is restored to its configured value. All static routes with the **ldp-sync** option enabled are also activated at the expiry of the timer.

Given the above behavior, it is recommended that the user configures the sync timer to a value which allows enough time for both the LDP IPv4 and LDP IPv6 sessions to come up.

## 5.23.11 BFD Operation

The operation of BFD over a LDP interface tracks the next-hop of prefix IPv4 and prefix IPv6 in addition to tracking of the LDP peer address of the Hello adjacency over that link. This tracking is required as LDP can now resolve both IPv4 and IPv6 prefix FECs over a single IPv4 or IPv6 LDP session and, as such, the next-hop of a prefix will not necessarily match the LDP peer source address of the Hello adjacency. The failure of either or both of the BFD session tracking the FEC next-hop and the one tracking the Hello adjacency will cause the LFA backup NHLFE for the FEC to be activated, or the FEC to be re-resolved if there is no FRR backup.

The following CLI command allows the user to decide if they want to track only with an IPv4 BFD session, only with an IPv6 BFD session, or both:

```
config>router>ldp>if-params>if>bfd-enable [ipv4] [ipv6]
```

This command provides the flexibility required in case the user does not need to track both Hello adjacency and next-hops of FECs. For example, if the user configures **bfd-enable ipv6** only to save on the number of BFD sessions, then LDP will track the IPv6 Hello adjacency and the next-hops of IPv6 prefix FECs. LDP will 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 will be black-holed. If the BFD tracking the IPv6 Hello adjacency times out, then all IPv4 and IPv6 prefix FECs will be updated.

The tracking of a mLDP FEC has the following behavior:

- IPv4 and IPv6 mLDP FECs will only be tracked with the Hello adjacency because they do not have the concept of downstream next-hop.
- The upstream LSR peer for an mLDP FEC supports the multicast upstream FRR procedures, and the upstream peer will be tracked using the Hello adjacency on each link or the IPv6 transport address if there is a T-LDP session.
- The tracking of a targeted LDP peer with BFD does not change with the support of IPv6 peers. BFD tracks the transport address conveyed by the Hello adjacency which bootstrapped the LDP IPv6 session.

## 5.23.12 Services Using SDP with an LDP IPv6 FEC

The SDP of type **LDP** with **far-end** and **tunnel-farend** options using IPv6 addresses is supported. The addresses need not be of the same family (IPv6 or IPv4) for the SDP configuration to be allowed. The user can have an SDP with an IPv4 (or IPv6) control plane for the T-LDP session and an IPv6 (or IPv4) LDP FEC as the tunnel.

Because IPv6 LSP is only supported with LDP, the use of a **far-end** IPv6 address will not be allowed with a BGP or RSVP/MPLS LSP. In addition, the CLI will 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 within 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 which use LDP control plane (such as T-LDP VPLS and R-VPLS, VLL, and IES/VP RN spoke interface) will have the spoke-SDP (PW) signaled with an IPv6 T-LDP session when the **far-end** option is configured to an IPv6 address. The spoke-SDP for these services binds by default to an SDP that uses a LDP IPv6 FEC, which prefix matches the far end address. The spoke-SDP can use a different LDP IPv6 FEC or a LDP IPv4 FEC as the tunnel by configuring the **tunnel-far-end** option. In addition, the IPv6 PW control word is supported with both data plane packets and VCCV OAM packets. Hash label is also supported with the 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.

## 5.23.13 Mirror Services and Lawful Intercept

The user can configure a spoke-SDP bound to an LDP IPv6 LSP to forward mirrored packets from a mirror source to a remote mirror destination. In the configuration of the mirror destination service at the destination node, the remote-source command must use a spoke-SDP with a VC-ID that matches the one that is configured in the mirror destination service at the mirror source node. The **far-end** option will not be supported with an IPv6 address.

This also applies to the configuration of the mirror destination for a LI source.

### 5.23.13.1 Configuration at mirror source node

Use the following rules and syntax to configure at the mirror source node.

- The *sdp-id* must match an SDP which uses LDP IPv6 FEC
- Configuring *egress-vc-label* is optional.

config mirror mirror-dest 10

**CLI Syntax:**

```
no spoke-sdp sdp-id:vc-id
spoke-sdp sdp-id:vc-id [create]
 egress
 vc-label egress-vc-label
```

### 5.23.13.2 Configuration at mirror destination node

Use the following rules and syntax to configure 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 a LDP IPv6 SDP coming from mirror source node.
  - In the **spoke-sdp sdp-id:vc-id** command, *vc-id* should match that of the **spoke-sdp** configured in the mirror-destination context at mirror source node.
  - Configuring *ingress-vc-label* is optional; both static and t-ldp are supported.
- configure mirror mirror-dest 10 remote-source

**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 and LI will also be supported with PW redundancy feature when the endpoint spoke-SDP, including the ICB, is using a LDP IPv6 tunnel.

## 5.23.14 Static Route Resolution to a LDP IPv6 FEC

An LDP IPv6 FEC can be used to resolve a static IPv6 route with an indirect next-hop matching the FEC prefix. The user configures a resolution filter to specify the LDP tunnel type to be selected from TTM:

**config>router>static-route-entry ip-prefix/prefix-length [mcast]**

```

indirect ip-address
 tunnel-next-hop
 [no] disallow-igp
 resolution {any | disabled | filter}
 resolution-filter
 [no] ldp

```

A static route of an IPv6 prefix cannot be resolved to an indirect next-hop using a LDP IPv4 FEC. An IPv6 prefix can only be resolved to an IPv4 next-hop using the 6-over-4 encapsulation by which the outer IPv4 header uses system IPv4 address as source and the next-hop as a destination. So the following example will return an error:

```

A:SRU4>config>router# static-route 3ffe::30/128 indirect 110.20.1.1 tunnel-next-hop
resolution-filter ldp

```

MINOR: CLI LDP not allowed for 6over4.

## 5.23.15 IGP Route Resolution to a LDP IPv6 FEC

LDP IPv6 shortcut for IGP IPv6 prefix is supported. The following commands allow a user to select if shortcuts must be enabled for IPv4 prefixes only, for IPv6 prefixes only, or for both.

**config>router>ldp-shortcut [ipv4][ipv6]**

```
ldp-shortcut [ipv4] [ipv6]
no ldp-shortcut
```

This CLI command has the following behaviors:

- When executing a pre-Release 13.0 config file, the existing command is converted as follows:  
**config>router>ldp-shortcut** changed to **config>router>ldp-shortcut ipv4**
- If the user enters the command without the optional arguments in the Release 13.0 CLI, it defaults to enabling shortcuts for IPv4 IGP prefixes:  
**config>router>ldp-shortcut** changed to **config>router>ldp-shortcut ipv4**
- When the user enters both IPv4 and IPv6 arguments in the Release 13.0 CLI, shortcuts for both IPv4 and IPv6 prefixes are enabled:  
**config>router>ldp-shortcut ipv4 ipv6**

## 5.23.16 OAM Support with LDP IPv6

MPLS OAM tools **lsp-ping** and **lsp-trace** are updated to operate with LDP IPv6 and support the following:

- use of IPv6 addresses in the echo request and echo reply messages, including in 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 is updated to support both LDP IPv4 and IPv6 target FEC stack TLVs. Specifically:

1. The IP family (IPv4/IPv6) of the UDP/IP echo request message will always match the family of the LDP target FEC stack TLV as entered by the user in the **prefix** option.
2. 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 will be used. If the user entered a source IP address of a different family than the LDP target FEC stack TLV, an error is returned and the test command is aborted.
3. The IP family of the UDP/IP echo reply message must match that of the received echo request message.
4. For **lsp-trace**, the downstream information in DSMAP/DDMAP will be encoded as the same family as the LDP control plane of the link LDP or targeted LDP session to the downstream peer.
5. The sender node inserts the experimental value of 65503 in the Router Alert Option in the echo request packet's IPv6 header as per RFC 5350. Once a value is allocated by IANA for MPLS OAM as part of *draft-ietf-mpls-oam-ipv6-rao*, it will be updated.

Finally, **vccv-ping** and **vccv-trace** for a single-hop PW are updated to support IPv6 PW FEC 128 and FEC 129 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. The value of the Channel Type field is set to 0x57 indicates that the Associated Channel carries an IPv6 packet, as per RFC 4385.

## 5.23.17 LDP IPv6 Interoperability Considerations

### 5.23.17.1 Interoperability with Implementations Compliant with draft-ietf-mpls-ldp-ipv6

SR OS implementation uses a 128-bit LSR-ID as defined in *draft-pdutta-mpls-ldp-v2* to establish an LDP IPv6 session with a peer LSR. This is done such 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, as is the common practice with LDP IPv4 in existing customer deployments. More importantly, this allows for the establishment of control plane independent LDP IPv4 and LDP IPv6 sessions between two LSRs over the same interface or set of interfaces. The SR OS implementation allows for two separate LDP IPv4 and LDP IPv6 sessions between two LSRs over the same interface or a set of interfaces because each session uses a unique LSR-ID (32-bit for IPv4 and 128-bit for IPv6).

The SR OS LDP implementation does not interoperate with an implementation using a 32-bit LSR-ID as defined in *draft-ietf-mpls-ldp-ipv6* to establish an IPv6 LDP session. The latter specifies an LSR can send both IPv4 and IPv6 Hellos over an interface such that it can establish either an IPv4 or an IPv6 LDP session with LSRs on the same subnet. It thus does not allow for separate LDP IPv4 and LDP IPv6 LDP sessions between two routers.

The SR OS LDP implementation should interoperate with an implementation using a 32-bit LSR-ID as defined in *draft-ietf-mpls-ldp-ipv6* to establish an IPv4 LDP session and to resolve both IPv4 and IPv6 prefix FECs.

The SR OS LDP implementation otherwise complies with all other aspects of *draft-ietf-mpls-ldp-ipv6*, including the support of the dual-stack capability TLV in the Hello message. The latter is used by an LSR to inform its peer that it is capable of establishing either an LDP IPv4 or LDP IPv6 session and to convey the IP family preference for the LDP Hello adjacency and 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 SR OS implementation has a separate session for each IP family between two LSRs and, as such, this TLV is used to indicate the family preference and to also indicate that it supports resolving IPv6 FECs over an IPv4 LDP session.

### **5.23.17.2 Interoperability with Implementations Compliant with RFC 5036 for IPv4 LDP Control Plane Only**


This implementation 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*. When introducing an LSR based on the SR OS in a LAN with a broadcast interface, it can peer with third party LSR implementations which support *draft-ietf-mpls-ldp-ipv6* and LSRs which do not. When its peers using IPv4 LDP control plane with a third-party LSR implementation which does not support it, the advertisement of IPv6 addresses or IPv6 FECs to that peer may cause it to bring down the IPv4 LDP session.

In other words, there are deployed third-party LDP implementations which are compliant with RFC 5036 for LDP IPv4, but which are not compliant with RFC 5036 for handling IPv6 address or IPv6 FECs over an LDP IPv4 session. To address this issue, *draft-ietf-mpls-ldp-ipv6* modifies RFC 5036 by requiring implementations complying with *draft-ietf-mpls-ldp-ipv6* to check for the dual-stack capability TLV in the IPv4 Hello message from the peer. Without the peer advertising this TLV, an LSR must not send IPv6 addresses and FECs to that peer. SR OS implementation implements this change.

## 5.24 LDP Process Overview

Figure 79 displays the process to provision basic LDP parameters.

**Figure 79 LDP Configuration and Implementation**



(\*) if some of the needed resources are not available consider implementing stricter import-policies and/or enabling the per-peer fec-limit functionality.

MPLS\_01



## 5.25 LDP-IGP Synchronization

The SR OS supports the synchronization of an IGP and LDP based on a solution described in RFC 5443, which consists of setting the cost of a restored link to infinity to give both the IGP and LDP time to converge. When a link is restored after a failure, the IGP sets the link cost to infinity and advertises it. The actual value advertised in OSPF is 0xFFFF (65535). The actual value advertised in IS-IS regular metric is 0x3F (63) and in IS-IS wide-metric is 0xFFFFFE (16777214). This synchronization feature is not supported on RIP interfaces.

When the LDP synchronization timer subsequently expires, the actual cost is put back and the IGP will readvertise it and use it at the next SPF computation. The LDP synchronization timer is configured using the following command:

```
config>router>if> [no] ldp-sync-timer seconds
```

The SR OS also supports an LDP End of LIB message, as defined in RFC 5919, that allows a downstream node to indicate to its upstream peer that it has advertised its entire label information base. The effect of this on the IGP-LDP synchronization timer is described below.

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

Next, an LDP Hello adjacency is brought up with the neighbor. The LDP synchronization timer is started by the IGP when the LDP session to the neighbor is up over the interface. This is to allow time for the label-FEC bindings to be exchanged.

When the LDP synchronization timer expires, the link cost is restored and is readvertised. The IGP will announce a new best next hop and LDP will use it if the label binding for the neighbor's FEC is available.

If the user changes the cost of an interface, the new value is advertised at the next flooding of link attributes by the IGP. However, if the LDP synchronization timer is still running, the new cost value will only be advertised after the timer expires. The new cost value will also be advertised after the user executes any of the following commands:

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

- **config>router>ospf>disable-ldp-sync**
- **router>isis>disable-ldp-sync**

If the user changes the value of the LDP synchronization timer parameter, the new value will take effect at the next synchronization event. If the timer is still running, it will continue to use the previous value.

If parallel links exist to the same neighbor, then the bindings and services should remain up as long as there is one interface that is up. However, the user-configured LDP synchronization timer still applies on the interface that failed and was restored. In this case, the router will only consider this interface for forwarding after the IGP readvertises its actual cost value.

The LDP End of LIB message is used by a node to signal completion of label advertisements, using a FEC TLV with the Typed Wildcard FEC element for all negotiated FEC types. This is done even if the system has no label bindings to advertise. The SR OS also supports the Unrecognized Notification TLV (RFC 5919) that indicates to a peer node that it will ignore unrecognized status TLVs. This indicates to the peer node that it is safe to send End of LIB notifications even if the node is not configured to process them.

The behavior of a system that receives an End of LIB status notification is configured through the CLI on a per-interface basis:

**config>router>if>[no] ldp-sync-timer *seconds* end-of-lib**

If the **end-of lib** option is not configured, then the LDP synchronization timer is started when the LDP Hello adjacency comes up over the interface, as described above. Any received End of LIB LDP messages are ignored.

If the **end-of-lib** option is configured, then the system will behave as follows on the receive side:

- The **ldp-sync-timer** is started.
- If LDP End of LIB Typed Wildcard FEC messages are received for every FEC type negotiated for a given session to an LDP peer for that IGP interface, the **ldp-sync-timer** is terminated and processing proceeds as if the timer had expired, that is, by restoring the IGP link cost.
- If the **ldp-sync-timer** expires before the LDP End of LIB messages are received for every negotiated FEC type, then the system restores the IGP link cost.
- The receive side will drop any unexpected End of LIB messages.

If the **end-of-lib** option is configured, then the system will also send out an End of LIB message for prefix and P2MP FECs once all FECs are sent for all peers that have advertised the Unrecognized Notification Capability TLV.

See the SR OS *Router Configuration Guide* for the CLI command descriptions for LDP-IGP Synchronization.



---

## 5.26 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](#)

## 5.27 LDP Configuration Overview

When the 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 distribution path (SDP). When signaling is *off*, labels must be manually configured when the SDP is bound to a service.

## 5.28 Basic LDP Configuration

This chapter provides information to configure LDP and remove configuration examples of common configuration tasks.

The LDP protocol instance is created in the no shutdown (enabled) state.

The following displays the default LDP configuration.

```
A:ALA-1>config>router>ldp# info

session-parameters
exit
interface-parameters
exit
targeted-session
exit
no shutdown
```

```

A:ALA-1>config>router>ldp#
```

## 5.29 Common Configuration Tasks

This section provides information to configure:

- [Enabling LDP](#)
- [Configuring FEC Originate Parameters](#)
- [Configuring Graceful-Restart Helper Parameters](#)
- [Applying Export and Import Policies](#)
- [Targeted Session Parameters](#)
- [Interface Parameters](#)
- [Session Parameters](#)
- [Interface Parameters](#)

### 5.29.1 Enabling LDP

LDP must be enabled in order for the protocol to be active. MPLS does not need to be enabled on the router except if the network interface uses the Packet over Sonet (POS) encapsulation (Sonet path encapsulation type set to ppp-auto). In this case, MPLS must be enabled and the interface name added into MPLS to allow for the MPLSCP to come up on the PPP link between the two peers and for MPLS to be used on the interface. MPLS is enabled in the config>router>mpls context.

Use the following syntax to enable LDP on a router:

**CLI Syntax:**    ldp

**Example:**       config>router# ldp

The following displays the enabled LDP configuration.

```
A:ALA-1>config>router# info

...
#-----
echo "LDP Configuration"
#-----
 ldp
 session-parameters
```

```
 exit
 interface-parameters
 exit
 targeted-session
 exit
 exit

...
A:ALA-1>config>router#
```

## 5.29.2 Configuring FEC Originate Parameters

A FEC can be added to the LDP IP prefix database with a specific label operation on the node. Permitted operations are pop or swap. For a swap operation, an incoming label can be swapped with a label in the range of 16 to 1048575. If a swap-label is not configured then the default value is 3.

A route table entry is required for a FEC with a pop operation to be advertised. For a FEC with a swap operation, a route-table entry must exist and user configured next-hop for swap operation must match one of the next-hops in route-table entry.

Use the following syntax to configure FEC originate parameters:

**CLI Syntax:**

```
config>router>ldp
fec-originate ip-prefix/mask [advertised-label in-label]
next-hop ip-address [swap-label out-label]
fec-originate ip-prefix/mask [advertised-label in-label]
pop
```

The following displays a FEC originate configuration example.

```
A:ALA-5>config>router# info

 fec-originate 100.1.1.1/32 pop
 fec-originate 100.2.1.1/32 advertised-label 1000 next-hop 10.10.1.2
 fec-originate 100.3.1.1/32 advertised-label 1001 next-hop 10.10.2.3
swap-label 131071
session-parameters
exit
interface-parameters
exit
targeted-session
exit
 exit

A:ALA-5>config>router>ldp#
```

### 5.29.3 Configuring Graceful-Restart Helper Parameters

Graceful-restart helper advertises to its LDP neighbors by carrying the fault tolerant (FT) session TLV in the LDP initialization message, assisting the LDP in preserving its IP forwarding state across the restart. Nokia's recovery is self-contained and relies on information stored internally to self-heal. This feature is only used to help third-party routers without a self-healing capability to recover.

Maximum recovery time is the time (in seconds) the sender of the TLV would like the receiver to wait, after detecting the failure of LDP communication with the sender.

Neighbor liveness time is the time (in seconds) the LSR is willing to retain its MPLS forwarding state. The time should be long enough to allow the neighboring LSRs to re-sync all the LSPs in a graceful manner, without creating congestion in the LDP control plane.

Use the following syntax to configure graceful-restart parameters:

**CLI Syntax:**    `config>router>ldp`  
                  `[no] graceful-restart`

### 5.29.4 Applying Export and Import Policies

Both inbound and outbound label binding filtering are supported. Inbound filtering allows a route policy to control the label bindings an LSR accepts from its peers. An import policy can accept or reject label bindings received from LDP peers.

Label bindings can be filtered based on:

- Neighbor — Match on bindings received from the specified peer.
- Prefix-list — Match on bindings with the specified prefix/prefixes.

Outbound filtering allows a route policy to control the set of LDP label bindings advertised by the LSR. An export policy can control the set of LDP label bindings advertised by the router. By default, label bindings for only the system address are advertised and propagate all FECs that are received. All other local interface FECs can be advertised using policies.



**Note:** The system IP address and static FECs cannot be blocked using an export policy.



Matches can be based on:

- All — all local subnets.
- Match — match on bindings with the specified prefix/prefixes.

Use the following syntax to apply import and export policies:

**CLI Syntax:**

```
config>router>ldp
export policy-name [policy-name...(upto 32 max)]
import policy-name [policy-name...(upto 32 max)]
```

The following displays export and import policy configuration examples.

```
A:ALA-1>config>router# info

export "LDP-export"
fec-originate 100.1.1.1/32 pop
fec-originate 100.2.1.1/32 advertised-label 1000 next-hop 10.10.1.2
import "LDP-import"
session-parameters
exit
interface-parameters
exit
targeted-session
exit

A:ALA-1>config>router#
```

## 5.29.5 Targeted Session Parameters

Use the following syntax to specify **targeted-session** parameters:

**CLI Syntax:**

```
config>router# ldp
targeted-session
 disable-targeted-session
 export-prefixes policy-name [policy-name...(up to 5
 max)]
 ipv4
 hello timeout factor
 keepalive timeout factor
 import-prefixes policy-name [policy-name...(up to 5
 max)]
 peer ip-address
 hello timeout factor
 keepalive timeout factor
 no shutdown
 tunneling
 lsp lsp-name
```

The following example displays an LDP configuration example:

```
A:ALA-1>config>router>ldp# info

...
 targeted-session
 ipv4
 hello 120 3
 keepalive 120 3
 exit
 peer 10.10.10.104
 hello 240 3
 keepalive 240 3
 exit
 exit

A:ALA-1>config>router>ldp#
```

## 5.29.6 Interface Parameters

Use the following syntax to configure interface parameters:

**CLI Syntax:**

```
config>router# ldp
interface-parameters
 interface ip-int-name [dual stack]
 bfd-enable [ipv4] [ipv6]
 ipv4/ipv6
 hello timeout factor
 keepalive timeout factor
 transport-address {system | interface}
 no shutdown
 ipv4/ipv6
 hello timeout factor
 keepalive timeout factor
 transport-address {system | interface}
```

The following example displays an interface parameter configuration example:

```
A:ALA-1>config>router>ldp# info

...
 interface-parameters
 interface "to-DUT1" dual-stack
 ipv4
 hello 240 3
 keepalive 240 3
 exit
 exit
 exit

A:ALA-1>config>router>ldp#
```

## 5.29.7 Session Parameters

Use the following syntax to specify interface parameters:

**CLI Syntax:**

```
config>router# ldp
session-parameters
 peer ip-address
 auth-keychain name
 authentication-key [authentication-key | hash-
 key] [hash | hash2]
 ttl-security min-ttl-value [log log-id]
```

The following example displays an LDP configuration example:

```
A:ALA-1>config>router>ldp# info

export "LDP-export"
import "LDP-import"
session-parameters
 peer 1.1.1.1
 exit
 peer 10.10.10.104
 exit
exit
tcp-session-parameters
 peer-transport 10.10.10.104
 authentication-key "E7GtYNZHTAaQqVMRDbfNIZpLtHg4ECOk" hash2
 exit
exit
interface-parameters
 interface "to-DUT1" dual-stack
 ipv4
 hello 240 3
 keepalive 240 3
 exit
 exit
exit
targeted-session
 ipv4
 hello 120 3
 keepalive 120 3
 exit
 peer 10.10.10.104
 hello 240 3
 keepalive 240 3
 exit
exit
exit

A:ALA-1>config>router>ldp#
```

## 5.29.8 LDP Signaling and Services

When LDP is enabled, targeted sessions can be established to create remote adjacencies with nodes that are not directly connected. When service distribution paths (SDPs) are configured, extended discovery mechanisms enable LDP to send periodic targeted hello messages to the SDP far-end point. The exchange of LDP hellos trigger session establishment. The SDP signaling default enables **tldp**. The service SDP uses the targeted-session parameters configured in the **config>router>ldp>targeted-session** context.

The SDP LDP and LSP commands are mutually exclusive; either one LSP can be specified or LDP can be enabled. If LDP is already enabled on an MPLS SDP, then an LSP cannot be specified on the SDP. If an LSP is specified on an MPLS SDP, then LDP cannot be enabled on the 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, refer to the *SR OS Services Guide*.

The following example displays the command syntax usage to configure enable LDP on an MPLS SDP:

**CLI Syntax:**

```
config>service>sdp#
ldp
signaling {off | tldp}
```

The following displays an example of an SDP configuration showing the signaling default tldp enabled.

```
A:ALA-1>config>service>sdp# info detail

description "MPLS: to-99"
far-end 10.10.10.99
signaling tldp
path-mtu 4462
keep-alive
hello-time 10
hold-down-time 10
max-drop-count 3
timeout 5
no message-length
no shutdown
exit
no shutdown

A:ALA-1>config>service>sdp#
```

The following displays an example of an SDP configuration for the 7750 SR, showing the signaling default tldp enabled.

```
A:ALA-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

```

```
A:ALA-1>config>service>sdp#
```

The following shows a working configuration of LDP over RSVP-TE (1) where tunnels look like the second example (2):

### Example 1: LDP over RSVP-TE

```
*A:ALA-1>config>router>ldp# info
```

```

prefer-tunnel-in-tunnel
interface-parameters
 interface "port-1/1/3"
 exit
 interface "port-lag-1"
 exit
exit
targeted-session
 peer 10.51.0.1
 shutdown
 tunneling
 lsp "to_P_1"
 exit
 exit
 peer 10.51.0.17
 shutdown
 tunneling
 lsp "to_P_6"
 exit
 exit
exit
exit

```

```
*A:ALA-1>config>router>ldp#
```

### Example 2: Tunnels

```
ALA-1>config>router>if-attr# info
```

```

admin-group "lower" value 2
admin-group "upper" value 1

```

```

*A:ALA-1>config>router>mpls# info

 resignal-timer 30
 interface "system"
 exit
 interface "port-1/1/3"
 exit
 interface "port-lag-1"
 exit
 path "dyn"
 no shutdown
 exit
 lsp "to_P_1"
 to 10.51.0.1
 cspf
 fast-reroute facility
 exit
 primary "dyn"
 exit
 no shutdown
 exit
 lsp "to_P_6"
 to 10.51.0.17
 cspf
 fast-reroute facility
 exit
 primary "dyn"
 exit
 no shutdown
 exit
 no shutdown

*A:ALA-1>config>router>mpls#
```

## 5.30 LDP Configuration Management Tasks

This section discusses the following LDP configuration management tasks:

- [Disabling LDP](#)
- [Modifying Targeted Session Parameters](#)
- [Modifying Interface Parameters](#)
- [Modifying Interface Parameters](#)

## 5.30.1 Disabling LDP

The **no ldp** command disables the LDP protocol on the router. All parameters revert to the default settings. LDP must be shut down before it can be disabled.

Use the following command syntax to disable LDP:

**CLI Syntax:**   no ldp  
                  shutdown

## 5.30.2 Modifying Targeted Session Parameters

The modification of LDP targeted session parameters does not take effect until the next time the session goes down and is re-establishes. Individual parameters cannot be deleted. The no form of a **targeted-session** parameter command reverts modified values back to the default. Different default parameters can be configured for IPv4 and IPv6 LDP targeted hello adjacencies.

The following example displays the command syntax usage to revert targeted session parameters back to the default values:

**Example:**

```
config>router# ldp
config>router>ldp# targeted-session
config>router>ldp>tcp-session-params>peer# no
authentication-key
config>router>ldp>targ-session# no disable-targeted-
session
config>router>ldp>targ-session>ipv4# no hello
config>router>ldp>targ-session>ipv4# no keepalive
config>router>ldp>targ-session# no peer 10.10.10.104
```

The following output displays the default values:

```
A:ALA-1>config>router>ldp>targeted# info detail

no disable-targeted-session
no import-prefixes
no export-prefixes
ipv4
no hello
no keepalive
no hello-reduction
exit
ipv6
no hello
no keepalive
no hello-reduction
```

```
exit

A:ALA-1>config>router>ldp>targeted#
```

### 5.30.3 Modifying Interface Parameters

Individual parameters cannot be deleted. The **no** form of an **interface-parameter** command reverts modified values back to the defaults. The modification of LDP targeted session parameters does not take effect until the next time the session goes down and is re-establishes.

The following output displays the default values:

```
A:ALA-1>config>router>ldp>if-params>if# info detail

no bfd-enable
ipv4
 no hello
 no keepalive
 no local-lsr-id
 fec-type-capability
 p2mp-ipv4 enable
 exit
 no transport-address
 no shutdown
exit
no shutdown

```



## 5.31 LDP Command Reference

### 5.31.1 Command Hierarchies

#### 5.31.1.1 LDP Commands

```

config
 — router
 — [no] ldp
 — [no] aggregate-prefix-match
 — prefix-exclude policy-name [policy-name...(up to 5 max)]
 — no prefix-exclude
 — [no] shutdown
 — [no] class-forwarding
 — egress-statistics
 — [no] fec-prefix ip-prefix[/mask]
 — accounting-policy policy-id
 — no accounting-policy
 — [no] collect-stats
 — [no] shutdown
 — [no] entropy-label-capability
 — export policy-name [policy-name...(up to 5 max)]
 — no export
 — [no] export-tunnel-table policy-name [policy-name...(up to 5 max)]
 — export-tunnel-table policy-name [policy-name...(up to 5 max)]
 — fast-reroute [backup-sr-tunnel]
 — no fast-reroute
 — fec-originate ip-prefix/mask [advertised-label in-label] [swap-label out-label]
 interface interface-name
 — fec-originate ip-prefix/mask [advertised-label i in-label] next-hop ip-address
 [swap-label out-label]
 — fec-originate ip-prefix/mask [advertised-label in-label] next-hop ip-address
 [swap-label out-label] interface interface-name
 — fec-originate ip-prefix/mask [advertised-label in-label] pop
 — no fec-originate ip-prefix/mask interface interface-name
 — no fec-originate ip-prefix/mask next-hop ip-address
 — no fec-originate ip-prefix/mask next-hop ip-address interface interface-
 name
 — no fec-originate ip-prefix/mask pop
 — [no] graceful-restart
 — maximum-recovery-time interval
 — no maximum-recovery-time
 — neighbor-liveness-time interval
 — no neighbor-liveness-time
 — [no] implicit-null-label

```

---

```

— import policy-name [policy-name...(up to 5 max)]
— no import
— interface-parameters
 — [no] interface ip-int-name [dual-stack]
 — bfd-enable [ipv4][ipv6]
 — no bfd-enable
 — [no] shutdown
 — [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 | interface-name
 interface-name}
 — 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 | interface-name
 interface-name}
 — 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
— label-withdrawal-delay seconds

```

---

```

— no label-withdrawal-delay
— [no] legacy-ipv4-lsr-interop
— [no] mcast-upstream-frr
— mp-mbb-time interval
— no mp-mbb-time
— [no] prefer-tunnel-in-tunnel
— session-parameters
 — [no] peer ip-address
 — [no] adv-adj-addr-only
 — [no] dod-label-distribution
 — export-addresses policy-name [policy-name...(up to 5 max)]
 — no export-addresses
 — export-prefixes policy-name [policy-name...(up to 5 max)]
 — no export-prefixes
 — fec-limit limit [log-only] [threshold percentage]
 — no fec-limit
 — [no] fec129-cisco-interop
 — fec-type-capability
 — p2mp {enable | disable}
 — prefix-ipv4 {enable | disable}
 — prefix-ipv6 {enable | disable}
 — import-prefixes policy-name [policy-name...(up to 5 max)]
 — no import-prefixes
 — [no] pe-id-mac-flush-interop
— [no] shortcut-transit-ttl-propagate
— [no] shortcut-local-ttl-propagate
— [no] shutdown
— targeted-session
 — [no] disable-targeted-session
 — export-prefixes policy-name [policy-name...(up to 5 max)]
 — no export-prefixes
 — import-prefixes policy-name [policy-name...(up to 5 max)]
 — no import-prefixes
 — ipv4
 — hello timeout factor
 — no hello
 — hello-reduction {enable factor | disable}
 — no hello-reduction
 — keepalive timeout factor
 — no keepalive
 — ipv6
 — hello timeout factor
 — no hello
 — hello-reduction {enable factor | disable}
 — no hello-reduction
 — keepalive timeout factor
 — no keepalive
 — peer ip-address
 — no peer ip-address
 — [no] bfd-enable
 — hello timeout factor
 — no hello
 — hello-reduction {enable factor | disable}
 — no hello-reduction

```

- **keepalive** *timeout factor*
- **no keepalive**
- **local-lsr-id** *interface-name*
- **no local-lsr-id**
- **[no] shutdown**
- **[no] tunneling**
  - **[no] lsp**
- **[no] peer-template** *template-name*
  - **[no] bfd-enable**
  - **hello** *timeout factor*
  - **no hello**
  - **hello-reduction** {**enable** *factor* | **disable**}
  - **no hello-reduction**
  - **keepalive** *timeout factor*
  - **no keepalive**
  - **local-lsr-id** *interface-name*
  - **no local-lsr-id**
  - **[no] shutdown**
  - **[no] tunneling**
- **peer-template-map** *template-name* **policy** *peer-prefix-policy1* [*peer-prefix-policy2... up to 5*]
- **no peer-template-map** *template-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] path-mtu-discovery**
    - **ttl-security** *min-ttl-value*
    - **no ttl-security**
- **tunnel-down-damp-time** *seconds*
- **no tunnel-down-damp-time**
- **[no] ldp-shortcut**

## 5.31.2 Command Descriptions

### 5.31.2.1 Generic Commands

ldp

**Syntax**    **[no] ldp**

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command creates the context to configure an LDP parameters. LDP is not enabled by default and must be explicitly enabled (<b>no shutdown</b>).</p> <p>To suspend the LDP protocol, use the <b>shutdown</b> command. Configuration parameters are not affected.</p> <p>The <b>no</b> form of the command deletes the LDP protocol instance, removing all associated configuration parameters. The LDP instance must first be disabled with the <b>shutdown</b> command before being deleted.</p> |
| <b>Default</b>     | none (LDP must be explicitly enabled)                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## ldp-shortcut

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ldp-shortcut</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>router                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command enables the resolution of IGP routes using LDP LSP across all network interfaces participating in the IS-IS and OSPF routing protocol in the system.</p> <p>When LDP shortcut is enabled, LDP populates the routing table with next-hop entries corresponding to all prefixes for which it activated an LDP FEC. For a given prefix, two route entries are populated in the system routing table. One corresponds to the LDP shortcut next-hop and has an owner of LDP. The other one is the regular IP next-hop. The LDP shortcut next-hop always has preference over the regular IP next-hop for forwarding user packets and specified control packets over a given outgoing interface to the route next-hop.</p> <p>All user and specified control packets for which the longest prefix match in RTM yields the FEC prefix will be forwarded over the LDP LSP.</p> <p>When an IPv4 packet is received on an ingress network interface, a subscriber IES interface, or a regular IES interface, the lookup of the packet by the ingress forwarding engine will result in the packet being sent labeled with the label stack corresponding to the NHLFE of the LDP LSP when the preferred RTM entry corresponds to an LDP shortcut.</p> <p>If the preferred RTM entry corresponds to an IP next-hop, the IPv4 packet is forwarded without a label.</p> <p>When ECMP is enabled and multiple equal-cost next-hops exit for the IGP route, the ingress forwarding engine will spray the packets for this route based on hashing routine currently supported for IPv4 packets. When the preferred RTM entry corresponds to an LDP shortcut route, spraying will be performed across the multiple next-hops for the LDP FEC. The FEC next-hops can either be direct link LDP neighbors or T-LDP neighbors reachable over RSVP LSPs in the case of LDP-over-RSVP but not both.</p> <p>When the preferred RTM entry corresponds to a regular IP route, spraying will be performed across regular IP next-hops for the prefix.</p> |

The **no** form of this command disables the resolution of IGP routes using LDP shortcuts.

**Default** no ldp-shortcut

## shutdown

**Syntax** [no] shutdown

**Context** config>router>ldp  
config>router>ldp>if-params>if  
config>router>ldp>if-params>if>ipv4  
config>router>ldp>if-params>if>ipv6  
config>router>ldp>targ-session>peer  
config>router>ldp>targ-session>peer-template  
config>router>ldp>egr-stats>fec-prefix  
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 may be deleted. For an LDP interface, the **shutdown** command exists under the main interface context and under each of the interface IPv4 and IPv6 contexts.

- **shutdown** under the **interface** context brings down both IPv4 and IPv6 Hello adjacencies and stops Hello transmission in both contexts.
- **shutdown** under the **interface** IPv4 or IPv6 contexts brings down the Hello adjacency and stops Hello transmission in that context only.

The user can also delete the entire IPv4 or IPv6 context under the interface with the **no ipv4** or **no ipv6** command which in addition to bringing down the Hello adjacency will delete the configuration

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

## aggregate-prefix-match

**Syntax** [no] aggregate-prefix-match

**Context** config>router>ldp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>The command enables the use by LDP of the aggregate prefix match procedures.</p> <p>When this option is enabled, LDP performs the following procedures for all prefixes. When an LSR receives a FEC-label binding from an LDP neighbor for a given specific FEC1 element, it will install the binding in the LDP FIB if:</p> <ul style="list-style-type: none"> <li>• It is able to perform a successful longest IP match of the FEC prefix with an entry in the routing table, and</li> <li>• The advertising LDP neighbor is the next-hop to reach the FEC prefix.</li> </ul> <p>When such a FEC-label binding has been installed in the LDP FIB, then LDP programs an NHLFE entry in the egress data path to forward packets to FEC1. It also advertises a new FEC-label binding for FEC1 to all its LDP neighbors.</p> <p>When a new prefix appears in the routing table, LDP inspects the LDP FIB to determine if this prefix is a better match (a more specific match) for any of the installed FEC elements. For any FEC for which this is true, LDP may have to update the NHLFE entry for this FEC.</p> <p>When a prefix is removed from the routing table, LDP inspects the LDP FIB for all FEC elements which matched this prefix to determine if another match exists in the routing table. If so, it updates the NHLFE entry accordingly. If not, it sends a label withdraw message to its LDP neighbors to remove the binding.</p> <p>When the next hop for a routing prefix changes, LDP updates the LDP FIB entry for the FEC elements which matched this prefix. It also updates the NHLFE entry for these FEC elements accordingly.</p> <p>The <b>no</b> form of this command disables the use by LDP of the aggregate prefix procedures and deletes the configuration. LDP resumes performing exact prefix match for FEC elements.</p> |
| <b>Default</b>     | no aggregate-prefix-match                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## prefix-exclude

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>prefix-exclude</b> <i>policy-name</i> [<i>policy-name...</i>(up to 5 max)]</p> <p><b>no prefix-exclude</b></p>                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>ldp>aggregate-prefix-match                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command specifies the policy name containing the prefixes to be excluded from the aggregate prefix match procedures. In this case, LDP will perform an exact match of a specific FEC element prefix as opposed to a longest match of one or more LDP FEC element prefixes, against this prefix when it receives a FEC-label binding or when a change to this prefix occurs in the routing table.</p> <p>The <b>no</b> form of this command removes all policies from the configuration.</p> |
| <b>Default</b>     | no prefix-exclude                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

---

class-forwarding

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>class-forwarding</b><br><b>no class-forwarding</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>ldp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | This command enables Class-Based Forwarding (CBF) capability. When this command is enabled, LDP prefixes resolved to a set of ECMP tunnel next-hops will have their packets forwarded to the LSP which is configured to carry the forwarding class the packet was classified to. As a pre-requisite for that forwarding behavior to happen, the user must have enabled IGP shortcuts in the routing instance, enabled ecmp in the global routing instance and have enabled the advertisement of unicast prefix FECs on the Targeted LDP (T-LDP) session to the peer. The user must also have assigned forwarding classes to LSPs (see <b>config&gt;router&gt;mpls&gt;lsp&gt;class-forwarding</b> ). |

Class based forwarding will occur when a set of ECMP tunnel next hops is considered consistent from a CBF perspective. It is the case if at least one CBF configuration (**fc** or **default-lsp**) is assigned to one or more LSPs. If no LSP of the set has the **default-lsp** option assigned, one LSP will automatically be selected for that assignment by LDP (the one with the lowest tunnel-id within the set of LSPs with one or more forwarding classes assigned). Multiple LSPs can have a same forwarding class assigned. However, for each of these forwarding classes only a single LSP will be used to forward packets classified into this forwarding class. That LSP is the one with the lowest tunnel-id amongst those sharing a given forwarding class. Similarly, multiple LSPs can have the **default-lsp** configuration assigned. Only a single one will be designated to be the Default LSP. That LSP is the one with the lowest tunnel-id amongst those with the **default-lsp** configuration assigned.

Under normal conditions, LDP prefix packets will be sprayed over a set of ECMP tunnel next-hops by selecting either the LSP to which is assigned the forwarding class of the packets, if one exists, or the Default LSP, if one does not exist. However, If the IOM detects that the LSP to which is assigned a forwarding class is not usable, it will switch the forwarding of packets classified to that forwarding class into the Default LSP, and if the IOM detects that the Default LSP is not usable, then it will revert to regular ECMP spraying across all tunnels in the set of ECMP tunnel next-hops. In other words, the CBF is suspended until LDP downloads a new consistent set of tunnel next-hops for the FEC.

This command only applies to LSR forwarding LDP FEC prefix packets over a set of MPLS LSPs using IGP shortcuts. It does not apply to LER forwarding of shortcut packets over LDP FEC, which is resolved to a set of MPLS LSPs using IGP shortcuts, nor does it apply to LER forwarding of packets of VPRN and Layer-2 services, which use auto-binding to LDP when the LDP FEC is resolved to a set of MPLS LSPs using IGP shortcuts.

The **no** form of this command disables the class based forwarding capability. In that case, LDP prefixes resolved to a set of ECMP tunnel next-hops will have their packets forwarded according to ECMP spraying.

|                |                     |
|----------------|---------------------|
| <b>Default</b> | no class-forwarding |
|----------------|---------------------|



## egress-statistics

|                    |                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>egress-statistics</b>                                                                                                                                            |
| <b>Context</b>     | config>router>ldp                                                                                                                                                   |
| <b>Description</b> | This command provides the context for the user to enter the LDP FEC prefix for the purpose of enabling egress data path statistics at the ingress LER for this FEC. |
| <b>Default</b>     | none                                                                                                                                                                |

## fec-prefix

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] fec-prefix</b> <i>ip-prefix[/mask]</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>router>ldp>egr-stats                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command configures statistics in the egress data path at the ingress LER or LSR for an LDP FEC. The user must execute the <b>no shutdown</b> command for this command to effectively enable statistics. The egress data path counters will be updated for both originating and transit packets. Originating packets may be service packets or IP user and control packets forwarded over the LDP LSP when used as an IGP shortcut. Transit packets of the FEC which are label switched on this node.</p> <p>When ECMP is enabled and multiple paths exist for a FEC, the same set of counters are updated for each packet forwarded over any of the NHLFEs associated with this FEC and for as long as this FEC is active.</p> <p>The statistics can be enabled on prefix FECs imported from both LDP neighbors and T-LDP neighbors (LDP over RSVP). Only /32 FEC prefixes are accepted. Service FECs, i.e., FEC 128 and FEC 129 are not valid. LDP FEC egress statistics are not collected at the Penultimate-Popping Hop (PHP) node for a LDP FEC using an implicit null label.</p> <p>The <b>no</b> form of this command disables the statistics in the egress data path and removes the accounting policy association from the LDP FEC.</p> |
| <b>Default</b>     | none                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## accounting-policy

|                    |                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>accounting-policy</b> <i>acct-policy-id</i><br><b>no accounting-policy</b>                                                                                                         |
| <b>Context</b>     | config>router>ldp>egr-stats                                                                                                                                                           |
| <b>Description</b> | <p>This command associates an accounting policy to the MPLS instance.</p> <p>An accounting policy must be defined before it can be associated else an error message is generated.</p> |

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

|                   |                                                                                                                                       |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | none                                                                                                                                  |
| <b>Parameters</b> | <i>acct-policy-id</i> — Enter the accounting <i>policy-id</i> as configured in the <b>config&gt;log&gt;accounting-policy</b> context. |
| <b>Values</b>     | 1 to 99                                                                                                                               |

## collect-stats

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] collect-stats</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>router>ldp>egr-stats                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command enables accounting and statistical data collection. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.</p> <p>When the <b>no collect-stats</b> 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 billing file. If a subsequent <b>collect-stats</b> command is issued then the counters written to the billing file include all the traffic while the <b>no collect-stats</b> command was in effect.</p> |
| <b>Default</b>     | collect-stats                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## export

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>export</b> <i>policy-name</i> [ <i>policy-name</i> ...up to 5 max]<br><b>no export</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>ldp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command specifies the export route policies used to determine which routes are exported to LDP. Policies are configured in the <b>config&gt;router&gt;policy-options</b> context.</p> <p>If no export policy is specified, non-LDP routes will not be exported from the routing table manager to LDP. LDP-learned routes will be exported to LDP neighbors. Present implementation of export policy (outbound filtering) can be used “only” to add FECs for label propagation. The export policy does not control propagation of FECs that an LSR receives from its neighbors.</p> <p>If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of 5 policy names can be specified.</p> <p>The <b>no</b> form of the command removes all policies from the configuration.</p> |

---

|                   |                                                                                                                                                                                                                                                                                                                                                 |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no export — No export route policies specified.                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b> | <p><i>policy-name</i> — 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 (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p>The specified name(s) must already be defined.</p> |

## fast-reroute

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fast-reroute [backup-sr-tunnel]<br/>[no] fast-reroute</b>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>router>ldp                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command enables LDP Fast-Reroute (FRR) procedures. When enabled, LDP uses both the primary next-hop and LFA next-hop, when available, for resolving the next-hop of an LDP FEC against the corresponding prefix in the routing table. This will result in LDP programming a primary NHLFE and a backup NHLFE into the forwarding engine for each next-hop of a FEC prefix for the purpose of forwarding packets over the LDP FEC. |

When any of the following events occurs, LDP instructs in the fast path the forwarding engines to enable the backup NHLFE for each FEC next-hop impacted by this event:

- An LDP interface goes operationally down, or is admin shutdown.
- An LDP session to a peer went down as the result of the Hello or Keep-Alive timer expiring.
- The TCP connection used by a link LDP session to a peer went down, due say to next-hop tracking of the LDP transport address in RTM, which brings down the LDP session.
- A BFD session, enabled on a T-LDP session to a peer, times-out and as a result the link LDP session to the same peer and which uses the same TCP connection as the T-LDP session goes also down.
- A BFD session enabled on the LDP interface to a directly connected peer, times out and brings down the link LDP session to this peer.

The **tunnel-down-dump-time** option or the **label-withdrawal-delay** option, when enabled, does not cause the corresponding timer to be activated for a FEC as long as a backup NHLFE is still available.

Because LDP can detect the loss of a neighbor/next-hop independently, it is possible that it switches to the LFA next-hop while IGP is still using the primary next-hop. Also, when the interface for the previous primary next-hop is restored, IGP may re-converge before LDP completed the FEC exchange with it neighbor over that interface. This may cause LDP to de-program the LFA next-hop from the FEC and blackhole traffic. In order to avoid this situation, it is recommended to enable IGP-LDP synchronization on the LDP interface.

When the SPF computation determines there is more than one primary next-hop for a prefix, it will not program any LFA next-hop in RTM. Thus, the LDP FEC will resolve to the multiple primary next-hops that provide the required protection.

The **backup-sr-tunnel** option enables the use of SR tunnel, as a remote LFA backup tunnel next-hop by an LDP FEC.

As a pre-requisite, the user must enable the stitching of LDP and SR in the LDP-to-SR direction. That is because the LSR must perform the stitching of the LDP ILM to SR tunnel when the primary LDP next-hop of the FEC fails. Thus LDP must listen to SR tunnels programmed by the IGP in TTM but the mapping server feature is not required.

Assuming the following:

- the **backup-sr-tunnel** option is enabled in LDP
- the **{loopfree-alternate remote-lfa}** option is enabled in the IGP instance
- LDP was able to resolve the primary next-hop of the LDP FEC in RTM

If the IGP LFA SPF does not find a regular LFA backup next-hop for a prefix of an LDP FEC, it will run the remote LFA algorithm. If IGP finds a remote LFA tunnel next-hop, LDP programs the primary next-hop of the FEC using a LDP NHLFE and programs the LFA backup next-hop using a LDP NHLFE pointing to the SR tunnel endpoint. Note that the LDP packet is not “tunneled” over the SR tunnel. The LDP label is actually stitched to the segment routing label stack. LDP points both the LDP ILM and the LTN to the backup LDP NHLFE which itself uses the SR tunnel endpoint.

The behavior of the feature is thus similar to the LDP-to-SR stitching feature, except the behavior is augmented to allow the stitching of an LDP ILM/LTN to a SR tunnel also when the primary LDP next-hop of the FEC fails.

If the LDP FEC primary next-hop failed and LDP has pre-programmed a remote LFA next-hop with a LDP backup NHLFE pointing to SR tunnel, the LDP ILM/LTN switches to it. Note that if for some reason the failure impacted only the LDP tunnel primary next-hop but not the SR tunnel primary next-hop, the LDP backup NHLFE will effectively point to the primary next-hop of the SR tunnel and traffic of the LDP ILM/LTN will follow this path instead of the remote LFA next-hop of the SR tunnel until the latter is activated.

This feature is limited to IPv4 /32 prefixes in both LDP and SR.

The **no** form of the command disables the exporting of BGP tunnels and SR tunnels to LDP.

The **no** form of this command disables LDP FRR.

**Default** no fast-reroute

## export-tunnel-table

**Syntax** **[no] export-tunnel-table** *policy-name* [*policy-name...*(up to 5 max)]

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>ldp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command enables exporting of BGP label route and SR tunnels from the TTM into LDP for the purpose of stitching a LDP FEC to a BGP tunnel or to a SR tunnel for the same destination prefix.</p> <p>The user enables the stitching of routes between LDP and BGP by configuring separately tunnel table route export policies in both protocols and enabling the advertising of RFC 3107, <i>Carrying Label Information in BGP-4</i>, formatted labeled routes for prefixes learned from LDP FECs.</p> <p>The route export policy in BGP instructs BGP to listen to LDP route entries in the CPM Tunnel Table. If a /32 LDP FEC prefix matches an entry in the export policy, BGP originates a BGP labeled route, stitches it to the LDP FEC, and re-distributes the BGP labeled route to its iBGP neighbors.</p> <p>The user adds LDP FEC prefixes with the statement '<b>from protocol ldp</b>' in the configuration of the existing BGP export policy at the global level, the peer-group level, or at the peer level using the commands:</p> <ul style="list-style-type: none"> <li>• config&gt;router&gt;bgp&gt;export <i>policy-name</i></li> <li>• config&gt;router&gt;bgp&gt;group&gt;export <i>policy-name</i></li> <li>• config&gt;router&gt;bgp&gt;group&gt;neighbor&gt;export <i>policy-name</i></li> </ul> <p>To indicate to BGP to evaluate the entries with the <b>from protocol ldp</b> statement in the export policy when applied to a specific BGP neighbor, use commands:</p> <ul style="list-style-type: none"> <li>• config&gt;router&gt;bgp&gt;group&gt;neighbor&gt;family label-ipv4 and</li> <li>• config&gt;router&gt;bgp&gt;group&gt;neighbor&gt;advertise-ldp-prefix</li> </ul> <p>Without the latter, only core IPv4 routes learned from RTM are advertised as BGP labeled routes to the neighbor. No stitching of LDP FEC to the BGP labeled route will be performed for this neighbor even if the same prefix was learned from LDP.</p> <p>The tunnel table route export policy in LDP instructs LDP to listen to BGP route entries in the CPM Tunnel Table. If a /32 BGP labeled route matches a prefix entry in the export policy, LDP originates an LDP FEC for the prefix, stitches it to the BGP labeled route, and re-distributes the LDP FEC to its iBGP neighbors.</p> <p>The user can add BGP labeled route prefixes with the statement <b>from protocol bgp</b> in the configuration of the LDP tunnel table export policy. The <b>from protocol</b> statement has an effect only when the protocol value is <b>ldp</b>. Policy entries with protocol values of <b>rsvp</b>, <b>bgp</b>, or any value other than <b>ldp</b> are ignored at the time the policy is applied to LDP.</p> |

In the LDP-to-SR data path direction, LDP listens to SR tunnel entries in the TTM. The user can restrict the export to LDP of SR tunnels from a specific prefix list. The user can also restrict the export to a specific IGP instance by optionally specifying the instance ID in the “from protocol” statement. The statement has an effect only when the protocol value is **isis** or **bgp**. Policy entries with any other protocol value are ignored at the time the 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 will follow the TTM selection rules to select a tunnel to which it will stitch the LDP ILM:

1. LDP selects the tunnel from the lowest TTM preference protocol.
2. If two or more of IS-IS protocol instances and BGP protocol have the same preference, then LDP selects the protocol using the default TTM protocol preference.
3. Within the same IGP protocol, LDP selects the lowest instance ID.

Whenever an LDP FEC primary next-hop cannot be resolved using an RTM route and a SR tunnel of type SR-ISIS to the same destination prefix matches a prefix entry in the export policy, LDP programs an LDP ILM and stitches it to the SR node-SID tunnel endpoint. LDP also originates an FEC for the prefix and re-distributes it to its LDP peers. When an LDP FEC is stitched to a SR tunnel, packets forwarded will benefit from the protection of the LFA/remote LFA backup next-hop of the SR tunnel.

When resolving a FEC, LDP will prefer resolution in RTM over that in TTM when both resolutions are possible. In other words the swapping of the LDP ILM to a LDP NHLFE is preferred over stitching it to a SR tunnel endpoint.

It is recommended to enable the `bfd-enable` option on the interfaces in LDP, IGP instance, and BGP contexts to speed up the failure detection and the activation of the SR LFA/remote-LFA backup next-hop or the BGP backup depending of the stitching operation.

This feature is limited to IPv4 /32 prefixes in LDP, BGP and SR.

The **no** form of the command disables the exporting of BGP tunnels and SR tunnels to LDP.

**Default** no export-tunnel-table — no tunnel table export route policy is specified.

**Parameters** *policy-name* — The export-tunnel-table route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains spaces, use double quotes to delimit the start and end of the string. The specified name(s) must already be defined.

## fec-originate

**Syntax** **fec-originate** *ip-prefix/mask* [**advertised-label** *in-label*] [**swap-label** *out-label*] **interface** **interface-name**  
**fec-originate** *ip-prefix/mask* [**advertised-label** *in-label*] **next-hop** *ip-address* [**swap-label** *out-label*]  
**fec-originate** *ip-prefix/mask* [**advertised-label** *in-label*] **next-hop** *ip-address* [**swap-label** *out-label*] **interface** **interface-name**

**fec-originate** *ip-prefix/mask* [**advertised-label** *in-label*] **pop**  
**no fec-originate** *ip-prefix/mask* **interface** *interface-name*  
**no fec-originate** *ip-prefix/mask* **next-hop** *ip-address*  
**no fec-originate** *ip-prefix/mask* **next-hop** *ip-address* **interface** *interface-name*  
**no fec-originate** *ip-prefix/mask* **pop**

**Context** config>router>ldp

**Description** This command defines a way to originate a FEC (with a swap action) for which the LSR is not egress, or to originate a FEC (with a pop action) for which the LSR is egress.

**Parameters** *ip-prefix/mask* — Specify information for the specified IP prefix and mask length.

**Values**

<ip-address/mask> ipv4-prefix - a.b.c.d  
 ipv4-prefix-le - [0..32]  
 ipv6-prefix x:x:x:x:x:x:x (eight 16-bit pieces)  
 x:x:x:x:x:d.d.d.d  
 x - [0..FFFF]H  
 d - [0..255]D  
 ipv6-prefix-le - [0..128]

**next-hop** — Specify the IP address of the next hop of the prefix.

**advertised-label** — Specify the label advertised to the upstream peer. If not configured, then the label advertised should be from the label pool. If the configured static label is not available then the IP prefix is not advertised.

**out-label** — Specify the LSR to swap the label. If configured, then the LSR should swap the label with the configured swap-label. If not configured, then the default action is pop if the next-hop parameter is not defined.

The next-hop, advertised-label, swap-label parameters are all optional. If next-hop is configured but no swap label specified, it will be a swap with label 3, such as, pop and forward to the next-hop. If the next-hop and swap-label are configured, then it is a regular swap. If no parameters are specified, a pop and route is performed.

**Values** 16 to 1048575

**in-label** — Specifies the number of labels to send to the peer associated with this FEC.

**Values** 32 to 1023

**pop** — Specifies to pop the label and transmit without the label.

**interface** *interface-name* — Specifies the name of the interface the label for the originated FEC is swapped to. For an unnumbered interface, this parameter is mandatory since there is no address for the next-hop. For a numbered interface, it is optional.

---

## graceful-restart

|                    |                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] graceful-restart</b>                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>ldp                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command enables graceful restart helper.</p> <p>The <b>no</b> form of the command disables graceful restart.</p> <p>Graceful restart helper configuration changes, enable/disable, or change of a parameter will cause the LDP session to bounce.</p> |
| <b>Default</b>     | no graceful-restart (disabled) — Graceful-restart must be explicitly enabled.                                                                                                                                                                                 |

## implicit-null-label

|                    |                                                                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] implicit-null-label</b>                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>router>ldp                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command enables the use of the implicit null label. Use this command to signal the IMPLICIT NULL option for all LDP FECs for which this node is the egress LER.</p> <p>The <b>no</b> form of this command disables the signaling of the implicit null label.</p> |
| <b>Default</b>     | no implicit-null-label                                                                                                                                                                                                                                                   |

## maximum-recovery-time

|                    |                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>maximum-recovery-time</b> <i>interval</i><br><b>no maximum-recovery-time</b>                                                     |
| <b>Context</b>     | config>router>ldp>graceful-restart                                                                                                  |
| <b>Description</b> | <p>This command configures the local maximum recovery time.</p> <p>The <b>no</b> form of the command returns the default value.</p> |
| <b>Default</b>     | 120                                                                                                                                 |
| <b>Parameters</b>  | <i>interval</i> — Specifies the length of time in seconds.                                                                          |
| <b>Values</b>      | 15 to 1800                                                                                                                          |

## neighbor-liveness-time

|               |                                               |
|---------------|-----------------------------------------------|
| <b>Syntax</b> | <b>neighbor-liveness-time</b> <i>interval</i> |
|---------------|-----------------------------------------------|



## no neighbor-liveness-time

|                    |                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>ldp>graceful-restart                                                                                      |
| <b>Description</b> | This command configures the neighbor liveness time.<br><br>The <b>no</b> form of the command returns the default value. |
| <b>Default</b>     | 120                                                                                                                     |
| <b>Parameters</b>  | <i>interval</i> — Specifies the length of time in seconds.<br><br><b>Values</b> 5 to 300                                |

## import

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>import</b> <i>policy-name</i> [ <i>policy-name</i> ...up to 5 max]<br><b>no import</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>router>ldp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | This command configures import route policies to determine which label bindings (FECs) are accepted from LDP neighbors. Policies are configured in the <b>config&gt;router&gt;policy-options</b> context.<br><br>If no import policy is specified, LDP accepts all label bindings from configured LDP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.<br><br>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 will override the previous command. A maximum of five policy names can be specified.<br><br>The <b>no</b> form of the command removes all policies from the configuration. |
| <b>Default</b>     | no import — No import route policies specified.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <i>policy-name</i> — The import route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.<br><br>The specified name(s) must already be defined.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## label-withdrawal-delay

|               |                                                                                  |
|---------------|----------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>label-withdrawal-delay</b> <i>seconds</i><br><b>no label-withdrawal-delay</b> |
|---------------|----------------------------------------------------------------------------------|

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>ldp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command specifies configures the time interval (in s), LDP will delay for the withdrawal of FEC-label binding it distributed to its neighbors when FEC is de-activated. When the timer expires, LDP then sends a label withdrawal for the FEC to all its neighbors. This is applicable only to LDP IPv4 prefix FECs and is not applicable to pseudowires (service FECs).</p> <p>When there is an upper layer (user of LDP) which depends of LDP control plane for failover detection then label withdrawal delay and tunnel-down-damp-time options must be set to 0.</p> <p>An example is PW redundancy where the primary PW doesn't have its own fast failover detection mechanism and the node depends on LDP tunnel down event to activate the standby PW.</p> |
| <b>Default</b>     | no label-withdrawal-delay                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <p><i>seconds</i> — Specifies the time that LDP delays the withdrawal of FEC-label binding it distributed to its neighbors when FEC is de-activated.</p> <p><b>Values</b> 3 to 120</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## legacy-ipv4-lsr-interop

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] legacy-ipv4-lsr-interop</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>ldp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command provides for a global LDP knob to allow interoperability with legacy IPv4 LSR implementations which do not comply with the processing of Hello TLVs with the U-bit set. Specifically, this feature disables the following Hello TLVs:</p> <ol style="list-style-type: none"><li>1. The ALU proprietary Interface Info TLV (0x3E05) in the Hello message sent to the peer. This also results in the non-generation of the ALU proprietary Hello Adjacency Status TLV (0x3E06) since the Interface Info TLV is not sent. This is performed in SR OS releases 12 and higher.</li><li>2. The RFC 7552 standard dual-stack capability TLV (0x701) and the ALU proprietary Adjacency capability TLV (0x3E07) in SR OS releases 13 and higher.</li></ol> |

## mcast-upstream-frr

|                    |                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] mcast-upstream-frr</b>                                                                                                                                 |
| <b>Context</b>     | config>router>ldp                                                                                                                                              |
| <b>Description</b> | <p>When LDP programs the primary ILM record in the data path, it provides the IOM with the This command enables the mLDP fast upstream switchover feature.</p> |

When this command is enabled and LDP is resolving a mLDP FEC received from a downstream LSR, it checks if an ECMP next-hop or a LFA next-hop exist to the root LSR node. 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. LDP then sends the corresponding labels to both upstream LSR nodes. In 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 will then start accepting packets.

In order 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
```

In order to make use of the LFA next-hop, the user must enable LFA using the following commands:

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

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

Enabling IP FRR or LDP FRR features is not strictly required since LDP only needs to know where the alternate next-hop to the root LSR is to be able to send the Label Mapping message to program the backup ILM at the initial signaling of the tree. Thus enabling the LFA option is sufficient. If however, 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 loss of **\*\*LDP session\*\*** to the upstream peer to which the primary ILM label had been advertised. We strongly recommended to perform the following:

- Enable BFD on all LDP interfaces to upstream LSR nodes. When BFD detects the loss of the last adjacency to the upstream LSR, it will bring down immediately the LDP session which will cause the ILM to activate the backup ILM.
- If there is a concurrent TLDP adjacency to the same upstream LSR node, enable BFD on the T-LDP peer in addition to enabling it on the interface.
- 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 other reason than a failure of the interface or of the upstream LSR, routing and LDP will go out of sync. This means the backup ILM will remain activated until the next time SPF is rerun by IGP. By enabling IGP-LDP synchronization feature, the advertised link metric will be changed to max value as soon as the LDP session goes down. This in turn will trigger an SPF and LDP will likely download a new set of primary and backup ILMs.

The **no** form of this command disables the fast upstream switchover for mLDP FECs.

**Default**    no mcast-upstream-frr

---

tunnel-down-damp-time

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>tunnel-down-damp-time</b> <i>seconds</i><br><b>no tunnel-down-damp-time</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>router>ldp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command specifies the time interval (in s), that LDP waits before posting a tunnel down event to the Tunnel Table Manager (TTM).</p> <p>When LDP can no longer resolve a FEC and de-activates it, it de-programs the NHLFE in the data path. It will however delay deleting the LDP tunnel entry in the TTM until the tunnel-down-damp-time timer expires. This means users of the LDP tunnel, such as SDPs (all services) and BGP (L3 VPN), will not be notified immediately. Traffic is still blackholed because the forwarding engine NHLFE has been de-programmed.</p> <p>If the FEC gets resolved before the tunnel-down-damp-time timer expires, then LDP programs the forwarding engine with the new NHLFE and performs a tunnel modify event in TTM updating the dampened entry in 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 TTM which deletes the LDP tunnel.</p> <p>When there is an upper layer (user of LDP) which depends of LDP control plane for failover detection then label withdrawal delay and tunnel-down-damp-time options must be set to 0.</p> <p>An example is pseudowire redundancy where the primary PW doesn't have its own fast failover detection mechanism and the node depends on LDP tunnel down event to activate the standby PW.</p> <p>The <b>no</b> form of this command resumes the default value of this command.</p> |
| <b>Parameters</b>  | <p><i>seconds</i> — Specifies the time interval (in s), that LDP waits before posting a tunnel down event to the Tunnel Table Manager.</p> <p><b>Default</b>      3 seconds</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## keepalive

|                |                                                                                                                                                                                                                                                                                                                         |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>keepalive</b> <i>timeout factor</i><br><b>no keepalive</b>                                                                                                                                                                                                                                                           |
| <b>Context</b> | config>router>ldp>if-params>if>ipv4<br>config>router>ldp>if-params>if>ipv6<br>config>router>ldp>if-params>ipv4<br>config>router>ldp>if-params>ipv6<br>config>router>ldp>targ-session>ipv4<br>config>router>ldp>targ-session>ipv6<br>config>router>ldp>targ-session>peer<br>config>router>ldp>targ-session>peer-template |

**Description** This command configures the time interval (in s), that LDP waits before tearing down the session. The **factor** parameter derives the keepalive interval.

The **config>router>ldp>if-params>ipv6>keepalive** and **config>router>ldp>targ-session>ipv6>keepalive** commands are not supported on the 7450 ESS.

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.

When LDP session is being set up, the keepalive timeout is negotiated to the lower of the two peers. Once a 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 interface-parameters and targeted-session levels sets the **keepalive timeout** and the **keepalive factor** to the default value.

The **no** form of the command, at the interface level, sets the **keepalive timeout** and the **keepalive factor** to the value defined under the **interface-parameters** level.

The **no** form of the command, at the peer level, will set the **keepalive timeout** and the **keepalive factor** to the value defined under the **targeted-session** level.

The session must be flapped for the new settings to operate.

**Default** [Table 42](#) lists the default values.

**Table 42** Timeout Factor Defaults

| Context                             | Timeout                                            | Factor |
|-------------------------------------|----------------------------------------------------|--------|
| config>router>ldp>if-params         | 30                                                 | 3      |
| config>router>ldp>targ-session      | 40                                                 | 4      |
| config>router>ldp>if-params>if      | Inherits values from interface-parameters context. |        |
| config>router>ldp>targ-session>peer | Inherits values from targeted-session context.     |        |

**Parameters** *timeout* — 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

---

local-lsr-id

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>local-lsr-id</b> { <b>system</b>   <b>interface</b>   <b>interface-name</b> <i>interface-name</i> }<br><b>no local-lsr-id</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>router>ldp>if-params>if>ipv4/<br>config>router>ldp>if-params>if>ipv6<br>config>router>ldp>targ-session>peer<br>config>router>ldp>targ-session>peer-template                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command enables the use of the address of the local LDP interface, or any other network interface configured on the system, as the LSR-ID to establish link LDP Hello adjacency and LDP session with directly connected LDP peers. The network interface can be a loopback or not.</p> <p>Link LDP sessions to all peers discovered over a given LDP interface share the same local LSR-ID. However, LDP sessions on different LDP interfaces can use different network interface addresses as their local LSR-ID.</p> <p>By default, the LDP session to a peer uses the system interface address as the LSR-ID unless explicitly configured using the above command. The system interface must always be configured on the router, or the LDP protocol will not come up on the node. There is no requirement to include it in any routing protocol.</p> <p>At initial configuration, the LDP session to a peer will remain down while the network interface used as LSR-ID is down. LDP will not try to bring it up using the system interface.</p> <p>At any time the network IP interface used as LSR-ID goes down, the LDP sessions to all discovered peers using this LSR-ID go down.</p> <p>If the user changes the LSR-ID value on the fly between <b>system</b>, <b>interface</b>, and <i>interface-name</i> while the LDP session is up, LDP will immediately tear down all sessions using this LSR-ID and will attempt to re-establish them using the new LSR-ID.</p> <p>When an interface other than system is used as the LSR-ID, the transport connection (TCP) for the link LDP session will also use the address of that interface as the transport address. If <b>system</b> or <b>interface</b> value is configured in the <b>config&gt;router&gt;ldp&gt;if-params&gt;if&gt;ipv4/ipv6&gt;transport-address</b> context, it will be overridden.</p> <p>The <b>no</b> form of the command returns to the default behavior in which case the system interface address is used as the LSR-ID.</p> |
| <b>Default</b>     | no local-lsr-id                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <i>interface-name</i> — Specifies the name, up to 32 character 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 (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## tunneling

|                    |                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] tunneling</b>                                                                               |
| <b>Context</b>     | config>router>ldp>targ-session>peer<br>config>router>ldp>targ-session>peer-template                 |
| <b>Description</b> | This command enables LDP over tunnels.<br><br>The <b>no</b> form of the command disables tunneling. |
| <b>Default</b>     | no tunneling                                                                                        |

## lsp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] lsp <i>lsp-name</i></b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>router>ldp>targ-session>tunneling                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | This command configures a specific LSP destined to this peer and to be used for tunneling of LDP FEC over RSVP. A maximum of 4 RSVP LSPs can be explicitly used for tunneling LDP FECs to the T-LDP peer.<br><br>It is not necessary to specify any RSVP LSP in this context unless there is a need to restrict the tunneling to selected LSPs. All RSVP LSPs with a to address matching that of the T-LDP peer are eligible by default. The user can also exclude specific LSP names by using the <b>ldp-over-rsvp exclude</b> command in the <b>configure-&gt;router-&gt;mpls-&gt;lsp <i>lsp-name</i></b> context. |

### 5.31.2.2 Interface Parameters Commands

## interface-parameters

|                    |                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>interface-parameters</b>                                                                                                                                                                                                                             |
| <b>Context</b>     | config>router>ldp                                                                                                                                                                                                                                       |
| <b>Description</b> | This command enables the context to configure LDP interfaces and parameters applied to LDP interfaces. The user can configure different default parameters for IPv4 and IPv6 LDP interfaces by entering <b>ipv4</b> or <b>ipv6</b> as the next command. |

## ipv4

|               |             |
|---------------|-------------|
| <b>Syntax</b> | <b>ipv4</b> |
|---------------|-------------|

**Context** config>router>ldp>interface parameters

**Description** This command enables the context to configure LDP interfaces and parameters applied to an IPv4 LDP interface.

## ipv6

**Syntax** **ipv6**

**Context** config>router>ldp>interface parameters

**Description** This command enables the context to configure LDP interfaces and parameters applied to an IPv6 LDP interface.

This command is not supported on the 7450 ESS.

## hello

**Syntax** **hello** *timeout factor*  
**no hello**

**Context** config>router>ldp>if-params>if>ipv4  
config>router>ldp>if-params>if>ipv6  
config>router>ldp>if-params>ipv4  
config>router>ldp>if-params>ipv6  
config>router>ldp>targ-session>ipv4  
config>router>ldp>targ-session>ipv6  
config>router>ldp>targ-session>peer  
config>router>ldp>targ-session>peer-template

**Description** This command configures the time interval to wait before declaring a neighbor down. The **factor** parameter derives the hello interval.

The **config>router>ldp>if-params>ipv6>hello** and **config>router>ldp>targ-session>ipv6>hello** commands are not supported on the 7450 ESS.

Hold time is local to the system and 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.

When LDP session is being set up, the holddown time is negotiated to the lower of the two peers. Once a 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 level sets the **hello timeout** and the **hello factor** to the default values.

The **no** form of the command, at the interface level, will set the **hello timeout** and the **hello factor** to the value defined under the interface-parameters level.



The **no** form of the command, at the peer level, will set the **hello timeout** and the **hello factor** to the value defined under the targeted-session level.

The session needs to be flapped for the new settings to operate.

**Default** Table 43 lists the default values.

**Table 43 Hello Timeout Factors**

| Context                             | Timeout                                            | Factor |
|-------------------------------------|----------------------------------------------------|--------|
| config>router>ldp>if-params         | 15                                                 | 3      |
| config>router>ldp>targ-session      | 45                                                 | 3      |
| config>router>ldp>if-params>if      | Inherits values from interface-parameters context. |        |
| config>router>ldp>targ-session>peer | Inherits values from targeted-session context.     |        |

**Parameters** *timeout* — Configures the time interval (in s), that LDP waits before 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

## hello-reduction

**Syntax** **hello-reduction {enable factor | disable}**  
**no hello-reduction**

**Context** config>router>ldp>targ-session>ipv4  
config>router>ldp>targ-session>ipv6  
config>router>ldp>targ-session>peer  
config>router>ldp>targ-session>peer-template

**Description** This command enables the suppression of periodic targeted Hello messages between LDP peers once the targeted LDP session is brought up.

The **config>router>ldp>targ-session>ipv6>hello-reduction** command is not supported on the 7450 ESS.

When this feature is enabled, the target Hello adjacency is brought up by advertising the Hold-Time value the user configured in the “**hello timeout**” parameter for the targeted session. The LSR node will then start advertising an exponentially increasing Hold-Time value in the Hello message as soon as the targeted LDP session to the peer is up. Each new incremented Hold-Time value is sent in a number of Hello messages equal to the value of the

argument *factor*, which represents the dampening factor, before the next exponential value is advertised. This provides time for the two peers to settle on the new value. When the Hold-Time reaches the maximum value of 0xffff (binary 65535), the two peers will send Hello messages at a frequency of every  $[(65535-1)/\text{local helloFactor}]$  seconds for the lifetime of the targeted-LDP session (for example, if the local Hello Factor is three (3), then Hello messages will be sent every 21844 seconds).

The LSR node continues to compute the frequency of sending the Hello messages based on the minimum of its local Hold-time value and the one advertised by its peer as in RFC 5036. Thus for the targeted LDP session to suppress the periodic Hello messages, both peers must bring their advertised Hold-Time to the maximum value. If one of the LDP peers does not, the frequency of the Hello messages sent by both peers will continue to be governed by the smaller of the two Hold-Time values.

When the user enables the hello reduction option on the LSR node while the targeted LDP session to the peer is operationally up, the change will take effect immediately. In other words, the LSR node will start advertising an exponentially increasing Hold-Time value in the Hello message, starting with the current configured Hold-Time value.

When the user disables the hello reduction option while the targeted LDP session to the peer is operationally up, the change in the Hold-Time from 0xffff (binary 65535) to the user configured value for this peer will take effect immediately. The local LSR will immediately advertise the value of the user configured Hold-Time value and will not wait until the next scheduled time to send a Hello to make sure the peer adjusts its local hold timeout value immediately.

In general, any configuration change to the parameters of the T-LDP Hello adjacency (modifying the hello adjacency Hello Timeout or factor, enabling/disabling hello reduction, or modifying hello reduction factor) will cause the LSR node to trigger immediately an updated Hello message with the updated Hold Time value without waiting for the next scheduled time to send a Hello.

The **no** form of this command disables the hello reduction feature.

|                   |                                                                                            |
|-------------------|--------------------------------------------------------------------------------------------|
| <b>Default</b>    | no hello-reduction                                                                         |
| <b>Parameters</b> | <i>factor</i> — Specifies the integer that specifies the Hello reduction dampening factor. |
| <b>Values</b>     | 3 to 20                                                                                    |

interface

|                    |                                                                |
|--------------------|----------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] interface</b> <i>ip-int-name</i> [ <b>dual-stack</b> ] |
| <b>Context</b>     | config>router>ldp>interface-parameters                         |
| <b>Description</b> | 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.

The user can configure different parameters for IPv4 and IPv6 LDP interfaces by entering **ipv4** or **ipv6** as the next command.

- Parameters**
- ip-int-name* — The name of an existing interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.
- dual-stack* — This optional keyword allows the user to distinguish between configuration execs prior to SR OS Release 13.0 from those in Release 13.0, as the interface node implementation has changed in Release 13.0 to include new IPv4 and IPv6 contexts. The following are some of the key points for this keyword:
- If the keyword is provided, then IPv4 interface context will not be created. If it is not provided, the IPv4 interface context will be created. This will take care of execs of prior to Release 13.0 configurations on a router running SR OS Release 13.0.
  - This new keyword will always show in a Release 13.0 configuration.
  - When entering an already configured interface, there is no need to provide the keyword, but it will be ignored if provided.
  - When deleting a configured interface, the keyword will not be accepted in the **no** version of the **interface** command.

## ipv4

- Syntax** **[no] ipv4**
- Context** config>router>ldp>interface parameters>interface
- Description** This command enables the context to configure IPv4 LDP parameters applied to the interface.

## ipv6

- Syntax** **[no] ipv6**
- Context** config>router>ldp>interface parameters>interface
- Description** This command enables the context to configure IPv6 LDP parameters applied to the interface.

This command is not supported on the 7450 ESS.

## bfd-enable

- Syntax** **bfd-enable [ipv4][ipv6]**  
**no bfd-enable**

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>router>ldp>if-params>if                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command enables tracking of the Hello adjacency to an LDP peer using BFD.</p> <p>The <b>ipv6</b> option for this command is not supported on the 7450 ESS.</p> <p>When this command is enabled on an LDP interface, LDP registers with BFD and starts tracking the LSR-id of all peers it formed Hello adjacencies with over that LDP interface. The LDP hello mechanism is used to determine the remote address to be used for the BFD session. The parameters used for the BFD session, that is, transmit-interval, receive-interval, and multiplier are those configured under the IP interface in existing implementation:</p> <p><b>config&gt;router&gt;if&gt;bfd</b></p> <p>The operation of BFD over an LDP interface 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 since LDP can resolve both IPv4 and IPv6 prefix FECs over a single IPv4 or IPv6 LDP session and as such the next-hop of a prefix will not necessarily match the LDP peer source address of the Hello adjacency.</p> <p>The failure of either or both of the BFD session tracking the FEC next-hop and the one tracking the Hello adjacency will cause the LFA backup NHLFE for the FEC to be activated or the FEC to be re-resolved if there is no FRR backup.</p> <p>When multiple links exist to the same LDP peer, a Hello adjacency is established over each link and a separate BFD session is enabled on each LDP interface. If a BFD session times out on a specific link, LDP will immediately associate the LDP session with one of the remaining Hello adjacencies and trigger the LDP FRR procedures. As soon as the last Hello adjacency goes down due to BFD timing out, the LDP session goes down and the LDP FRR procedures will be triggered.</p> <p>The <b>no</b> form of this command disables BFD on the LDP interface.</p> |
| <b>Default</b>     | no bfd-enable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

## transport-address

|                    |                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>transport-address {interface   system}</b><br><b>no transport-address</b>                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>router>ldp>if-params>if>ipv4<br>config>router>ldp>if-params>if>ipv6<br>config>router>ldp>if-params>ipv4<br>config>router>ldp>if-params>ipv6                                                                                                                                                                      |
| <b>Description</b> | <p>This command configures the transport address to be used when setting up the LDP TCP sessions. The transport address can be configured as <b>interface</b> or <b>system</b>. The transport address can be configured globally (applies to all LDP interfaces) or per interface. The most specific value is used.</p> |

The **config>router>ldp>if-params>ipv6>transport-address** command is not supported on the 7450 ESS.

With the transport-address command, you can set up the LDP interface to the connection which can be set to the interface address or the system address. However, there can be an issue of which address to use when there are parallel adjacencies. This situation can not only happen with parallel links, it could be a link and a targeted adjacency since 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 for **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 match the adjacency to the session.

For any ILDP interface, as the **local-lsr-id** parameters is changed to **interface**, the **transport-address** configuration loses effectiveness. Since it will be ignored and the ILDP session will *always* use the relevant interface IP address as transport-address even though system is chosen.

The **no** form of the command, at the global level, sets the transport address to the default value.

The **no** form of the command, at the interface level, sets the transport address to the value defined under the global level.

**Default** system — The system IP address is used.

**Parameters** **interface** — 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, since only one LDP session is set up between two neighbors.

**system** — The system IP address is used to set up the LDP session between neighbors.

## mp-mbb-time

**Syntax** [no] **mp-mbb-time**

**Context** config>router>ldp

**Description** This command configures the maximum time a P2MP transit/bud node must wait before switching over to the new path if the new node does not send MBB TLV to inform of the availability of data plane.

The **no** form of command should configure the default timer of 3 seconds.

**Default** 3 seconds

---

|                   |                            |
|-------------------|----------------------------|
| <b>Parameters</b> | <i>interval</i> — seconds. |
| <b>Values</b>     | 1 to 10 seconds            |

### 5.31.2.3 Session Parameters Commands

#### session-parameters

|                    |                                                                         |
|--------------------|-------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>session-parameters</b>                                               |
| <b>Context</b>     | config>router>ldp                                                       |
| <b>Description</b> | This command enables the context to configure peer specific parameters. |

#### peer

|                    |                                                                                       |
|--------------------|---------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] <b>peer</b> <i>ip-address</i>                                                    |
| <b>Context</b>     | config>router>ldp>session-parameters                                                  |
| <b>Description</b> | This command configures parameters for an LDP peer.                                   |
| <b>Default</b>     | none                                                                                  |
| <b>Parameters</b>  | <i>ip-addr</i> — The IPv4 or IPv6 address of the LDP peer in dotted decimal notation. |

#### adv-adj-addr-only

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] <b>adv-adj-addr-only</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>router>ldp>session-params>peer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command provides a means for an LDP router to advertise only the local IPv4 or IPv6 interfaces it uses to establish hello adjacencies with an LDP peer. By default, when a router establishes an LDP session with a peer, it advertises in an LDP Address message the addresses of all local interfaces to allow the peer to resolve LDP FECs distributed by this router. Similarly, a router sends a Withdraw Address message to of all its peers to withdraw a local address if the corresponding interface went down or was deleted.</p> <p>This new option reduces CPU processing when a large number of LDP neighbors come up or go down. The new CLI option is strongly recommended in mobile backhaul networks where the number of LDP peers can be very large.</p> <p>The <b>no</b> form of this command reverts LDP to the default behavior of advertising all local interfaces.</p> |

## dod-label-distribution

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] dod-label-distribution</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>router>ldp>session-params>peer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command enables the use of the LDP Downstream-on-Demand (DoD) label distribution procedures.</p> <p>When this option is enabled, LDP will set the A-bit in the Label Initialization message when the LDP session to the peer is established. When both peers set the A-bit, they will both use the DoD label distribution method over the LDP session (RFC 5036).</p> <p>This feature can only be enabled on a link-level LDP session and therefore will apply to prefix labels only, not service labels.</p> <p>As soon as the link LDP session comes up, the router will send a label request to its DoD peer for the FEC prefix corresponding to the peer's LSR-id. The DoD peer LSR-id is found in the basic Hello discovery messages the peer used to establish the Hello adjacency with the router.</p> <p>Similarly if the router and the directly attached DoD peer entered into extended discovery and established a targeted LDP session, the router will immediately send a label request for the FEC prefix corresponding to the peer's LSR-id found in the extended discovery messages.</p> <p>However, the router will not advertise any &lt;FEC, label&gt; bindings, including the FEC of its own LSR-id, unless the DoD peer requested it using a Label Request Message.</p> <p>When the DoD peer sends a label request for any FEC prefix, the router will reply with a &lt;FEC, label&gt; binding for that prefix if the FEC was already activated on the router. If not, the router replies with a notification message containing the status code of "no route." The router will not attempt in the latter case to send a label request to the next-hop for the FEC prefix when the LDP session to this next-hop uses the DoD label distribution mode. Hence the reference to single-hop LDP DoD procedures.</p> <p>As soon as the link LDP session comes up, the router will send a label request to its DoD peer for the FEC prefix corresponding to the peer's LSR-id. The DoD peer LSR-id is found in the basic Hello discovery messages the peer used to establish the Hello adjacency with the router.</p> <p>Similarly if the router and the directly attached DoD peer entered into extended discovery and established a targeted LDP session, the router will immediately send a label request for the FEC prefix corresponding to the peer's LSR-id found in the extended discovery messages. Peer address has to be the peer LSR-ID address.</p> <p>The <b>no</b> form of this command disables the DoD label distribution with an LDP neighbor.</p> |
| <b>Default</b>     | no dod-label-distribution                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

---

## export-addresses

|                    |                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>export-addresses</b> <i>policy-name</i> [ <i>policy-name...</i> (up to 5 max)]<br><b>no export-addresses</b>                                                                                                                                                                                                                     |
| <b>Context</b>     | config>router>ldp>session-params>peer                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command specifies the export prefix policy to local addresses advertised to this peer.</p> <p>Policies are configured in the <b>config&gt;router&gt;policy-options</b> context. A maximum of five policy names can be specified.</p> <p>The <b>no</b> form of the command removes the policy from the configuration.</p>    |
| <b>Parameters</b>  | <i>policy-name</i> — The export-prefix route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains spaces, use double quotes to delimit the start and end of the string. The specified name(s) must already be defined. |

## export-prefixes

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] export-prefixes</b> <i>policy-name</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>router>ldp>session-params>peer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command specifies the export route policy used to determine which prefixes received from other LDP and T-LDP peers are re-distributed to this LDP peer via the LDP/T-LDP session to this peer. A prefix that is filtered out (deny) will not be exported. A prefix that is filtered in (accept) will be exported.</p> <p>If no export policy is specified, all FEC prefixes learned will be exported to this LDP peer. This policy is applied in addition to the global LDP policy and targeted session policy.</p> <p>Policies are configured in the <b>config&gt;router&gt;policy-options</b> context. A maximum of five policy names can be specified. Peer address has to be the peer LSR-ID address.</p> <p>The <b>no</b> form of the command removes the policy from the configuration.</p> |
| <b>Default</b>     | no export-prefixes - no export route policy is specified                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <i>policy-name</i> — The export-prefix route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains spaces, use double quotes to delimit the start and end of the string. The specified name(s) must already be defined.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

## fec-limit

|               |                                                                                   |
|---------------|-----------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>fec-limit</b> <i>limit</i> [ <i>log-only</i> ] [ <i>threshold percentage</i> ] |
|---------------|-----------------------------------------------------------------------------------|



## no fec-limit

**Context** config>router>ldp>session-params>peer

**Description** This command configures a limit on the number of FECs which an LSR will accept from a given peer and add into the LDP label database. The limit applies to the aggregate count of all FEC types including service FEC. Once the limit is reached, any FEC received will be released back to the peer. This behavior is different from the per-peer import policy which will still accept the FEC into the label database but will not resolve it.

When the FEC limit for a peer is reached, the LSR performs the following actions:

1. Generates a trap and a syslog message.
2. Generates a LDP notification message with the LSR overload status TLV, for each LDP FEC type including service FEC, to this peer only if this peer advertised support for the LSR overload sub-TLV via the LSR Overload Protection Capability TLV at session initialization.
3. Releases, with LDP Status Code of "No\_Label\_Resources", any new FEC, including service FEC, from this peer which exceeds the limit.

If a legitimate FEC is released back to a peer, while the FEC limit was exceeded, the user must have a means to replay that FEC back to the router LSR once the condition clears. This is done automatically if the peer is an SR-OS-based router and supports the LDP overload status TLV (SR OS 11.0R5 and higher). Third-party peer implementations must support the LDP overload status TLV or provide a manual command to replay the FEC.

The **threshold percent** option allows to set a threshold value when a trap and an syslog message are generated as a warning to the user in addition to when the limit is reached. The default value for the threshold when not configured is 90%.

The **log-only** option causes a trap and syslog message to be generated when reaching the threshold and limit. However, LDP labels are not released back to the peer.

If the user decreases the limit value such that it is lower than the current number of FECs accepted from the peer, the LDP LSR raises the trap for exceeding the limit. In addition, it will set overload for peers which signaled support for LDP overload protection capability TLV. However, no existing resolved FECs from the peer which does not support the overload protection capability TLV should be de-programmed or released.

A different trap is released when crossing the threshold in the upward direction, when reaching the FEC limit, and when crossing the threshold in the downward direction. However the same trap will not be generated more often than 2 minutes apart if the number of FECs oscillates around the threshold or the FEC limit.

**Default** no fec-limit

**Parameters** *limit* — Specify the aggregate count of FECs of all types which can be accepted from this LDP peer.

**log-only** — Specify only a trap and syslog message are generated when reaching the threshold and limit. However, LDP labels are not released back to the peer.

---

**threshold percent** — Specify the threshold value (as a percentage) that triggers a warning syslog message and trap to be sent.

## fec129-cisco-interop

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] fec129-cisco-interop</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>ldp>session-params>peer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command specifies whether LDP will provide translation between non-compliant FEC 129 formats of Cisco. Peer LDP sessions must be manually configured towards the non-compliant Cisco PEs.</p> <p>When enabled, Cisco non-compliant format will be used to send and interpret received label release messages i.e. the FEC129 SAll and TAll fields will be reversed.</p> <p>When the disabled, Cisco non-compliant format will not be used or supported. Peer address has to be the peer LSR-ID address.</p> <p>The <b>no</b> form of the command returns the default.</p> |
| <b>Default</b>     | no fec129-cisco-interop                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

## fec-type-capability

|                    |                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fec-type-capability</b>                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>ldp>session-params>peer<br>config>router>ldp>if-params>if>ipv4<br>config>router>ldp>if-params>if>ipv6                                                                                                                                         |
| <b>Description</b> | <p>This command enables or disables the advertisement of a FEC type on a given LDP session or Hello adjacency to a peer.</p> <p>The config&gt;router&gt;ldp&gt;if-params&gt;if&gt;ipv6&gt;fec-type-capability command is not supported on the 7450 ESS.</p> |

## p2mp

|                    |                                                                       |
|--------------------|-----------------------------------------------------------------------|
| <b>Syntax</b>      | <b>p2mp {enable   disable}</b>                                        |
| <b>Context</b>     | config>router>ldp>session-params>peer>fec-type-capability             |
| <b>Description</b> | This command enables or disables P2MP FEC capability for the session. |

## p2mp-ipv4

|                    |                                                                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>p2mp {enable   disable}</b>                                                                                                                                                                     |
| <b>Context</b>     | config>router>ldp>if-params>if>ipv4>fec-type-capability<br>config>router>ldp>if-params>if>ipv6>fec-type-capability<br>config>router>ldp>session-params>peer>fec-type-capability                    |
| <b>Description</b> | This command enables or disables IPv4 P2MP FEC capability on the interface.<br><br>The config>router>ldp>if-params>if>ipv6>fec-type-capability>p2mp-ipv4 command is not supported on the 7450 ESS. |

## p2mp-ipv6

|                    |                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>p2mp {enable   disable}</b>                                                                                                                                                  |
| <b>Context</b>     | config>router>ldp>if-params>if>ipv4>fec-type-capability<br>config>router>ldp>if-params>if>ipv6>fec-type-capability<br>config>router>ldp>session-params>peer>fec-type-capability |
| <b>Description</b> | This command enables or disables IPv6 P2MP FEC capability on the interface.<br><br>This command is not supported on the 7450 ESS.                                               |

## prefix-ipv4

|                    |                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>prefix-ipv4 {enable   disable}</b>                                                                                                                                                                             |
| <b>Context</b>     | config>router>ldp>if-params>if>ipv4>fec-type-capability<br>config>router>ldp>if-params>if>ipv6>fec-type-capability<br>config>router>ldp>session-params>peer>fec-type-capability                                   |
| <b>Description</b> | This command enables or disables IPv4 prefix FEC capability on the session or interface.<br><br>The config>router>ldp>if-params>if>ipv6>fec-type-capability>prefix-ipv4 command is not supported on the 7450 ESS. |

## prefix-ipv6

|                    |                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>prefix-ipv6 {enable   disable}</b>                                                                                                                                           |
| <b>Context</b>     | config>router>ldp>if-params>if>ipv4>fec-type-capability<br>config>router>ldp>if-params>if>ipv6>fec-type-capability<br>config>router>ldp>session-params>peer>fec-type-capability |
| <b>Description</b> | This command enables or disables IPv6 prefix FEC capability on the session or interface.                                                                                        |

---

This command is not supported on the 7450 ESS.

## import-prefixes

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] import-prefixes</b> <i>policy-name</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>router>ldp>session-params>peer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command configures the import FEC prefix policy to determine which prefixes received from this LDP peer are imported and installed by LDP on this node. If resolved these FEC prefixes are then re-distributed to other LDP and T-LDP peers. A FEC prefix that is filtered out (deny) will not be imported. A FEC prefix that is filtered in (accept) will be imported.</p> <p>If no import policy is specified, the node will import all prefixes received from this LDP/T-LDP peer. This policy is applied in addition to the global LDP policy and targeted session policy.</p> <p>Policies are configured in the <b>config&gt;router&gt;policy-options</b> context. A maximum of five policy names can be specified. Peer address has to be the peer LSR-ID address.</p> <p>The <b>no</b> form of the command removes the policy from the configuration.</p> |
| <b>Default</b>     | no import-prefixes - no import route policy is specified                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <i>policy-name</i> — The import-prefix route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains spaces, use double quotes to delimit the start and end of the string. The specified name(s) must already be defined.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## path-mtu-discovery

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>path-mtu-discovery</b><br><b>no path-mtu-discovery</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>router>ldp>tcp-session-params>peer-transport                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command enables Path MTU discovery for the associated TCP connections. When enabled, the MTU for the associated TCP session is initially set to the egress interface MTU. The DF bit is also set so that if a router along the path of the TCP connection cannot handle a packet of a particular size without fragmenting, it sends back an ICMP message to set the path MTU for the given session to a lower value that can be forwarded without fragmenting. If one or more transport addresses used in the Hello adjacencies to the same peer LSR are different from the LSR-ID value, the user must add each of the transport addresses to the path MTU discovery configuration as a separate peer. This means when the TCP connection is bootstrapped by a given Hello adjacency, the path MTU discovery can operate over that specific TCP connection by using its specific transport address.</p> |
| <b>Default</b>     | no path-mtu-discovery                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## pe-id-mac-flush-interop

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] pe-id-mac-flush-interop</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>router>ldp>session-params>peer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | This command enables the addition of the PE-ID TLV in the LDP MAC withdrawal (mac-flush) message, under certain conditions, and modifies the mac-flush behavior for interoperability with other vendors that do not support the flush-all-from-me vendor-specific TLV. This flag can be enabled on a per LDP peer basis and allows the flush-all-from-me interoperability with other vendors. When the pe-id-mac-flush-interop flag is enabled for a given peer, the current mac-flush behavior is modified in terms of mac-flush generation, mac-flush propagation and behavior upon receiving a mac-flush. |

The mac-flush generation will be changed depending on the type of event and according to the following rules:

- Any all-from-me mac-flush event will trigger a mac-flush all-but-mine message (RFC 4762 compliant format) with the addition of a PE-ID TLV. The PE-ID TLV contains the IP address of the sending PE.
- Any all-but-mine mac-flush event will trigger a mac-flush all-but-mine message WITHOUT the addition of the PE-ID TLV, as long as the source spoke-sdp is not part of an end-point.
- Any all-but-mine mac-flush event will trigger a mac-flush all-but-mine message WITH the addition of the PE-ID TLV, if the source spoke-sdp is part of an end-point and the spoke-sdp goes from down/standby state to active state. In this case, the PE-ID TLV will contain the IP address of the PE to which the previous active spoke-sdp was connected to.

Any other case will follow the existing mac-flush procedures.

When the pe-id-mac-flush-interop flag is enabled for a given LDP peer, the mac-flush ingress processing is modified according to the following rules:

- Any received all-from-me mac-flush will follow the existing mac-flush all-from-me rules regardless of the existence of the PE-ID.
- Any received all-but-mine mac-flush will take into account the received PE-ID, i.e. all the mac addresses associated to the PE-ID will be flushed. If the PE-ID is not included, the mac addresses associated to the sending PE will be flushed.
- Any other case will follow the existing mac-flush procedures.

When a mac-flush message has to be propagated (for an ingress sdp-binding to an egress sdp-binding) and the pe-id-mac-flush-interop flag is enabled for the ingress and egress TLDP peers, the following behavior is observed:

- If the ingress and egress bindings are spoke-sdp, the PE will propagate the mac-flush message with its own PE-ID.

- If the ingress binding is an spoke-sdp and the egress binding a mesh-sdp, the PE will propagate the mac-flush message without modifying the PE-ID included in the PE-ID TLV.
- If the ingress binding is a mesh-sdp and the egress binding an spoke-sdp, the PE will propagate the mac-flush message with its own PE-ID.
- When ingress and egress bindings are mesh-sdp, the mac-flush message is never propagated. This is the behavior regardless of the pe-id-mac-flush-interop flag configuration.

The PE-ID TLV is never added when generating a mac-flush message on a B-VPLS if the send-bvpls-flush command is enabled in the I-VPLS. In the same way, no PE-ID is added when propagating mac-flush from a B-VPLS to a I-VPLS when the propagate-mac-flush-from-bvpls command is enabled. Mac-flush messages for peers within the same I-VPLS or within the same B-VPLS domain follow the procedures described above.

**Default** no pe-id-mac-flush-interop

## prefer-tunnel-in-tunnel

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] prefer-tunnel-in-tunnel</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>ldp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command specifies to use tunnel-in-tunnel over a simple LDP tunnel. Specifically, the user packets for LDP FECs learned over this targeted LDP session can be sent inside an RSVP LSP which terminates on the same egress router as the destination of the targeted LDP session. The user can specify an explicit list of RSVP LSP tunnels under the Targeted LDP session or LDP will perform a lookup in the Tunnel Table Manager (TTM) for the best RSVP LSP. In the former case, only the specified LSPs will be considered to tunnel LDP user packets. In the latter case, all LSPs available to the TTM and which terminate on the same egress router as this targeted LDP session will be considered. In both cases, the metric specified under the LSP configuration is used to control this selection.</p> <p>The lookup in the TTM will prefer a LDP tunnel over an LDP-over-RSVP tunnel if both are available. Also, the tunneling operates on the dataplane only. Control packets of this targeted LDP session are sent over the IGP path.</p> |

## shortcut-transit-ttl-propagate

|                    |                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shortcut-transit-ttl-propagate</b>                                                                                                                                                                                       |
| <b>Context</b>     | config>router>ldp<br>config>router>mpls                                                                                                                                                                                          |
| <b>Description</b> | <p>This command configures the TTL handling of transit packets for all LSP shortcuts originating on this ingress LER. It applies to all LDP or RSVP LSPs that are used to resolve static routes, BGP routes, and IGP routes.</p> |

The user can enable or disable the propagation of the TTL from the header of an IP packet into the header of the resulting MPLS packet independently for local and transit packets forwarded over an LSP shortcut.

By default, the feature propagates the TTL from the header of transit IP packets into the label stack of the resulting MPLS packets forwarded over the LSP shortcut. This is referred to as Uniform mode.

When the **no** form of the command is enabled, TTL propagation is disabled on all transit IP packets received on any IES interface and destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack. This is referred to as Pipe mode.

**Default** shortcut-transit-ttl-propagate

## shortcut-local-ttl-propagate

**Syntax** [no] shortcut-local-ttl-propagate

**Context** config>router>ldp  
config>router>mpls

**Description** This command configures the TTL handling of locally generated packets for all LSP shortcuts originating on this ingress LER. It applies to all LDP or RSVP LSPs that are used to resolve static routes, BGP routes, and IGP routes.

The user can enable or disable the propagation of the TTL from the header of an IP packet into the header of the resulting MPLS packet independently for local and transit packets forwarded over an LSP shortcut.

Local IP packets include ICMP Ping, traceroute, and OAM packets, that are destined to a route that is resolved to the LSP shortcut. Transit IP packets are all IP packets received on any IES interface and destined to a route that is resolved to the LSP shortcut

By default, the feature propagates the TTL from the header of locally generated IP packets into the label stack of the resulting MPLS packets forwarded over the LSP shortcut. This is referred to as Uniform mode.

When the **no** form of the above command is enabled, TTL propagation is disabled on all locally generated IP packets, including ICMP Ping, traceroute, and OAM packets, that are destined to a route that is resolved to the LSP shortcut. In this case, a TTL of 255 is programmed onto the pushed label stack. This is referred to as Pipe mode.

**Default** shortcut-local-ttl-propagate

---

### 5.31.2.4 Targeted Session Commands

#### targeted-session

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>targeted-session</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>router>ldp                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>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 user can configure different default parameters for IPv4 and IPv6 LDP targeted hello adjacencies.</p> <p>The discovery messages for an indirect LDP session are addressed to the specified peer and not to the multicast address.</p> |
| <b>Default</b>     | none                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

#### bfd-enable

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] bfd-enable</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Context</b>     | config>router>ldp>targ-session>peer<br>config>router>ldp>targ-session>peer-template                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>This command enables the bidirectional forwarding detection (BFD) session for the selected TLDP session. By enabling BFD for a selected targeted session, the state of that session is tied to the state of the underneath BFD session between the two nodes.</p> <p>The parameters used for the BFD are set via the BFD command under the IP interface.</p> <p>The <b>no</b> form of this command removes the TLDP session operational state binding to the central BFD session one.</p> |
| <b>Default</b>     | no bfd-enable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

#### disable-targeted-session

|                    |                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] disable-targeted-session</b>                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>router>ldp>targ-session                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command disables support for SDP triggered automatic generated 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.</p> |



The **no** form of the command enables the set up of any targeted sessions.

**Default** no disable-targeted-session

## peer

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

**Context** config>router>ldp>targeted-session

**Description** This command configures parameters for an LDP peer.no

**Default** none

**Parameters** *ip-address* — The IPv4 or IPv6 address of the LDP peer in dotted decimal notation.

## peer-template-map

**Syntax** **peer-template-map** *template-name* **policy** *peer-prefix-policy1* [*peer-prefix-policy2..up to 5*]  
**no peer-template-map peer-template** *template-name*

**Context** config>router>ldp>targeted-session

**Description** This command enables the automatic creation of a targeted Hello adjacency and LDP session to a discovered peer. The user configures a targeted session peer parameter template and binds it to a peer prefix policy.

Each application of a targeted session template to a given prefix in the prefix list will result in the establishment of a targeted Hello adjacency to an LDP peer using the template parameters as long as the prefix corresponds to a router-id for a node in the TE database. As a result of this, the user must enable the traffic-engineering option in ISIS or OSPF. The targeted Hello adjacency will either trigger a new LDP session or will be associated with an existing LDP session to that peer.

Up to 5 peer prefix policies can be associated with a single peer template at all times. Also, the user can associate multiple templates with the same or different peer prefix policies. Thus multiple templates can match with a given peer prefix. In all cases, the targeted session parameters applied to a given peer prefix are taken from the first created template by the user. This provides a more deterministic behavior regardless of the order in which the templates are associated with the prefix policies.

Each time the user executes the above command, with the same or different prefix policy associations, or the user changes a prefix policy associated with a targeted peer template, the system re-evaluates the prefix policy. The outcome of the re-evaluation will tell LDP if an existing targeted Hello adjacency needs to be torn down or if an existing targeted Hello adjacency needs to have its parameters updated on the fly.

If a /32 prefix is added to (removed from) or if a prefix range is expanded (shrunk) in a prefix list associated with a targeted peer template, the same prefix policy re-evaluation described above is performed.

The template comes up in the **no shutdown** state and as such it takes effect immediately. Once a template is in use, the user can change any of the parameters on the fly without shutting down the template. In this case, all targeted Hello adjacencies are updated.

The SR OS supports multiple ways of establishing a targeted Hello adjacency to a peer LSR:

- User configuration of the peer with the targeted session parameters inherited from the **config>router>ldp>targeted-session** in the top level context or explicitly configured for this peer in the **config>router>ldp>targ-session>peer** context and which overrides the top level parameters shared by all targeted peers. Let us refer to the top level configuration context as the global context. Some parameters only exist in the global context; their value will always be inherited by all targeted peers regardless of which event triggered it.
- User configuration of an SDP of any type to a peer with the signaling tldp option enabled (default configuration). In this case the targeted session parameter values are taken from the global context.
- User configuration of a (FEC 129) PW template binding in a BGP-VPLS service. In this case the targeted session parameter values are taken from the global context.
- User configuration of a (FEC 129 type II) PW template binding in a VLL service (dynamic multi-segment PW). In this case the target session parameter values are taken from the global context
- User configuration of a mapping of a targeted session peer parameter template to a prefix policy when the peer address exists in the TE database (this feature). In this case, the targeted session parameter values are taken from the template.

Since the above triggering events can occur simultaneously or in any arbitrary order, the LDP code implements a priority handling mechanism in order to decide which event overrides the active targeted session parameters. The overriding trigger will become the owner of the targeted adjacency to a given peer. The following is the priority order:

- Priority 1: manual configuration of session parameters
- Priority 2: mapping of targeted session template to prefix policy.
- Priority 3: manual configuration of SDP, PW template binding in BGP-AD VPLS and in FEC 129 VLL.

Any parameter value change to an active targeted Hello adjacency caused by any of the above triggering events is performed on the fly by having LDP immediately send a Hello message with the new parameters to the peer without waiting for the next scheduled time for the Hello message. This allows the peer to adjust its local state machine immediately and maintains both the Hello adjacency and the LDP session in UP state. The only exceptions are the following:

- The triggering event caused a change to the local-lsr-id parameter value. In this case, the Hello adjacency is brought down which will also cause the LDP session to be brought down if this is the last Hello adjacency associated with the session. A new Hello adjacency and LDP session will then get established to the peer using the new value of the local LSR ID.
- The triggering event caused the targeted peer shutdown option to be enabled. In this case, the Hello adjacency is brought down which will also cause the LDP session to be brought down if this is the last Hello adjacency associated with the session.

Finally, the value of any LDP parameter which is specific to the LDP/TCP session to a peer is inherited from the **config>router>ldp>session-params>peer** context. This includes MD5 authentication, LDP prefix per-peer policies, label distribution mode (DU or DOD), etc.

The no form of this command deletes the binding of the template to the peer prefix list and brings down all Hello adjacencies to the discovered LDP peers.

## peer-template

|                    |                                                                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] peer-template template-name</b>                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>router>ldp>targeted-session                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | This command creates a targeted session peer parameter template that can be referenced in the automatic creation of targeted Hello adjacency and LDP session to a discovered peer.<br><br>The <b>no</b> form of command deletes the peer template. A peer template cannot be deleted if it is bound to a peer prefix list. |
| <b>Parameters</b>  | <i>template-name</i> — Specifies the template name to identify targeted peer template. It must be 32 characters maximum.                                                                                                                                                                                                   |

## export-prefixes

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>export-prefixes policy-name [policy-name...(up to 5 max)]</b><br><b>no export-prefixes</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>router>ldp>targeted-session                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | This command specifies the export route policy used to determine which FEC prefix label bindings are exported from a targeted LDP session. A route that is filtered out (deny) will not be exported. A route that is filtered in (accept) will be exported.<br><br>If no export policy is specified, all bindings learned through a targeted LDP session will be exported to all targeted LDP peers. This policy is applied in addition to the global LDP policy.<br><br>Policies are configured in the <b>config&gt;router&gt;policy-options</b> context. A maximum of five policy names can be specified. |

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

**Parameters**    *policy-name* — The export policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## import-prefixes

**Syntax**        **import-prefixes** *policy-name* [*policy-name*...(up to 5 max)]  
**no import-prefixes**

**Context**        config>router>ldp>targeted-session

**Description**    This command configures the import route policy to determine which FEC prefix label bindings are accepted from targeted LDP neighbors into this node. A label binding that is filtered out (deny) will not be imported. A route that is filtered in (accept) will be imported.

If no import policy is specified, this node session will accept all bindings from configured targeted LDP neighbors. This policy is applied in addition to the global LDP policy.

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.

**Parameters**    *policy-name* — The import policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.

## ipv4

**Syntax**        **ipv4**

**Context**        config>router>ldp>targeted-session

**Description**    This command enables the context to configure parameters applied to targeted sessions to all IPv4 LDP peers.

## ipv6

**Syntax**        **ipv6**

**Context**        config>router>ldp>targeted-session

**Description** This command enables the context to configure parameters applied to targeted sessions to all IPv6 LDP peers.

This command is not supported on the 7450 ESS.

### 5.31.2.5 TCP Session Parameters Commands

#### tcp-session-parameters

**Syntax** **tcp-session-parameters**

**Context** config>router>ldp

**Description** This command enables the context to configure parameters applicable to TCP transport session of an LDP session to remote peer.

#### peer-transport

**Syntax** **peer-transport** *ip-address*  
**no peer transport**

**Context** config>router>ldp>tcp-session-parameters

**Description** This command configures the peer transport address, that is, the destination address of the TCP connection, and not the address corresponding to the LDP LSR-ID of the peer.

**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*

**Context** config>router>ldp>tcp-session-params>peer-transport

**Description** This command configures TCP authentication keychain to use for the session.

**Parameters** *name* — Specifies the name of the keychain to use for the specified TCP session or sessions. This keychain allows the rollover of authentication keys during the lifetime of a session up to 32 characters in length. Peer address has to be the TCP session transport address.

---

authentication-key

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>authentication-key</b> [ <i>authentication-key</i>   <i>hash-key</i> ] [ <i>hash</i>   <i>hash2</i> ]<br><b>no authentication-key</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>router>ldp>tcp-session-params>peer-transport                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command specifies the authentication key to be used between LDP peers before establishing sessions. Authentication uses the MD-5 message-based digest. Peer address has to be the TCP session transport address. If one or more transport addresses used in the Hello adjacencies to the same peer LSR are different from the LSR-ID value, the user must add each of the transport addresses to the authentication-key configuration as a separate peer. This means when the TCP connection is bootstrapped by a given Hello adjacency, the authentication can operate over that specific TCP connection by using its specific transport address.</p> <p>The <b>no</b> form of this command disables authentication.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Default</b>     | none                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><i>authentication-key</i> — The authentication key. The key can be any combination of ASCII characters up to 16 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").</p> <p><i>hash-key</i> — 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 (" ").</p> <p>This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.</p> <p><b>hash</b> — Specifies the key is entered in an encrypted form. If the <b>hash</b> or <b>hash2</b> parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> or <b>hash2</b> parameter specified</p> <p><b>hash2</b> — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the <b>hash2</b> encrypted variable cannot be copied and pasted. If the <b>hash</b> or <b>hash2</b> parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the <b>hash</b> or <b>hash2</b> parameter specified.</p> |

## path-mtu-discovery

|                |                                                           |
|----------------|-----------------------------------------------------------|
| <b>Syntax</b>  | <b>path-mtu-discovery</b><br><b>no path-mtu-discovery</b> |
| <b>Context</b> | config>router>ldp>tcp-session-params>peer-transport       |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command enables Path MTU discovery for the associated TCP connections. When enabled, the MTU for the associated TCP session is initially set to the egress interface MTU. The DF bit is also set so that if a router along the path of the TCP connection cannot handle a packet of a particular size without fragmenting, it sends back an ICMP message to set the path MTU for the given session to a lower value that can be forwarded without fragmenting.</p> <p>If one or more transport addresses used in the Hello adjacencies to the same peer LSR are different from the LSR-ID value, the user must add each of the transport addresses to the path MTU discovery configuration as a separate peer. This means when the TCP connection is bootstrapped by a given Hello adjacency, the path MTU discovery can operate over that specific TCP connection by using its specific transport address.</p> |
| <b>Default</b>     | no path-mtu-discovery                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

## ttl-security

|                    |                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ttl-security</b> <i>min-ttl-value</i><br><b>no ttl-security</b>                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>router>ldp>tcp-session-params>peer-transport                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command configures TTL security parameters for incoming packets. When the feature is enabled, BGP/LDP will accept incoming IP packets from a peer only if the TTL value in the packet is greater than or equal to the minimum TTL value configured for that peer. Peer address has to be the TCP session transport address.</p> <p>The <b>no</b> form of the command disables TTL security.</p> |
| <b>Default</b>     | no ttl-security                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <i>min-ttl-value</i> — Specify the minimum TTL value for an incoming packet.<br><b>Values</b> 1 to 255                                                                                                                                                                                                                                                                                                  |





## 5.32 Show, Clear, Debug, and Tools Command Reference

- [Show Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)
- [Tools Commands](#)

### 5.32.1 Command Hierarchies

#### 5.32.1.1 Show Commands

```

show
 — router
 — ldp
 — bindings
 — 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 p2mp-id identifier root ip-address [rd rd] [summary |
 detail] [egress-lsp tunnel-id] [egress-if port-id] inner-root ip-
 address
 — active p2mp p2mp-id identifier root ip-address [rd rd] [summary |
 detail] [egress-lsp tunnel-id] inner-root ip-address

```

- **active** **p2mp** **p2mp-id** *identifier* **root** *ip-address* [**rd** *rd*] [**summary** | **detail**] [*egress-nh ip-address*] **inner-root** *ip-address*
- **active** **p2mp** **p2mp-id** *identifier* [**rd** *rd*] [**family**] [**summary** | **detail**] [*innermost-root ip-address*]
- **active** **p2mp** **source** *ip-address* **group** *mcast-address* [**family**] [**summary** | **detail**] [*innermost-root ip-address*]
- **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** **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*]
- **detail** [**session** *ip-addr* [*label-space*]] [*family*]
- **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** **p2mp-id** *identifier* **root** *ip-address* [**session** *ip-addr* [*label-space*]] [**summary** | **detail**] **inner-root** *ip-address*
- **p2mp** **p2mp-id** *identifier* [**rd** *rd*] [**session** *ip-addr* [*label-space*]] [**family**] [**summary** | **detail**] [*innermost-root ip-address*]
- **p2mp** **p2mp-id** *identifier* **root** *ip-address* **rd** *rd* [**session** *ip-addr* [*label-space*]] [**summary** | **detail**] **inner-root** *ip-address*
- **p2mp** [**session** *ip-addr* [*label-space*]] [**family**] [**summary** | **detail**] [*opaque-type opaque-type*]
- **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* **root** *ip-address* [**session** *ip-addr* [*label-space*]] [**family**] [**summary** | **detail**] **inner-root** *ip-address*
- **p2mp** **source** *ip-address* **group** *mcast-address* [**session** *ip-addr* [*label-space*]] [**family**] [**summary** | **detail**] [*innermost-root ip-address*]
- **prefixes** [*family*] [**summary** | **detail**] [*egress-if port-id*]
- **prefixes** [*family*] [**summary** | **detail**] [*egress-lsp tunnel-id*]
- **prefixes** [*egress-nh ip-address*] [*family*] [**summary** | **detail**]

```

— prefixes prefix ip-prefix/ip-prefix-length [summary | detail] [egress-if
 port-id]
— prefixes prefix ip-prefix/ip-prefix-length [summary | detail] [egress-lsp
 tunnel-id]
— prefixes prefix ip-prefix/ip-prefix-length [egress-nh ip-address]
 [summary | detail]
— 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]]
— session [family] [summary | detail] ip-addr[label-space]
— summary [session ip-addr[label-space]] [ipv4 | ipv6]

show
— router
 — 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]]
 — discovery peer [ip-address] [state state] [detail | summary] [session ip-
 addr[label-space]]
 — fec-egress-stats [ip-prefix/mask]
 — fec-egress-stats [family]
 — fec-originate ip-prefix/mask [operation-type]
 — fec-originate [operation-type] [family]
 — interface [ip-int-name] [detail] [family]
 — interface resource-failures [family]
 — parameters
 — session [ip-addr[label-space]] [session-type] [state state] [summary | detail]
 — session [ip-addr[label-space]] local-addresses [sent | recv] [family]
 — session [ip-addr[label-space]][sent | recv] overload [fec-type fec-type]
 — session [sent | recv] overload [fec-type fec-type] [family]
 — session [ip-addr[label-space]] statistics [packet-type] [session-type]
 — session statistics [packet-type] [session-type] [family]
 — session [session-type] [state state] [summary | detail] [family]
 — session-parameters [family]
 — session-parameters peer-ip-address
 — statistics
 — statistics-summary [active] [family]
 — status
 — targ-peer [ip-address] [detail]
 — targ-peer [detail] family
 — targ-peer resource-failures [family]
 — targ-peer-template [peer-template]
 — targ-peer-template-map [template-name]
 — targ-peer-template-map [template-name] peers
 — tcp-session-parameters

```

- **tcp-session-parameters** [family]
- **tcp-session-parameters** [keychain *keychain*]
- **tcp-session-parameters** [transport-peer-ip-address]

See OS OAM and Diagnostics Guide for tools command descriptions, syntax, and usage information.

### 5.32.1.2 Clear Commands

```
clear
 — router
 — ldp
 — fec-egress-statistics [ip-prefix/mask]
 — instance [ipv4 | ipv6]
 — interface [ip-int-name] [statistics] [ipv4 | ipv6]
 — peer [ip-address] [statistics]
 — resource-failures
 — session ip-addr[label-space] [statistics]
 — session ip-addr[label-space] overload [fec-type p2mp | prefixes sub-type
 sub-type]
 — session ip-addr[label-space] overload [fec-type svc-fec128 | svc-fec129]
 — statistics
```

### 5.32.1.3 Debug Commands

```
[no] debug
 — router
 — [no] ldp
 — [no] interface interface-name family
 — [no] event
 — [no] messages
 — [no] packet [detail]
 — hello [detail]
 — no hello
 — peer ip-address
 — [no] event
 — [no] bindings
 — [no] messages
 — [no] packet
 — hello [detail]
 — no hello
 — init [detail]
 — no init
 — [no] keepalive
 — label [detail]
 — no label
```

### 5.32.1.4 Tools Commands

See the *OAM and Diagnostics Guide* for CLI description and syntax.

```
tools
 — dump
 — ldp-treetrace {prefix ip-prefix/mask | manual-prefix ip-prefix/mask}[path-
 destination ip-address] [trace-tree]
 — router
 — ldp
 — fec vc-type vc-type agi agi
 — fec p2mp-id identifier root ip-address
 — fec prefix ip-address[/mask]
 — fec root ip-address source ip-address group mcast-address inner-
 root ip-address
 — fec root ip-address source ip-address group mcast-address [rd rd]
 — fec vc-type vc-type vc-id vc-id
 — fec vc-type vc-type agi agi saii-type2 global-id:prefix:ac-id taii-type2
 global-id:prefix:ac-id
 — instance
 — interface ip-int-name
 — memory-usage
 — peer ip-address
 — session ip-addr[label-space] [connection | peer | adjacency]
 — sockets
 — timers [session ip-addr[label-space]]
 — static-route ldp-sync-status
 — perform
 — router
 — isis
 — ldp-sync-exit
 — run-manual-spf
 — ospf/ospf3
 — ldp-sync-exit
 — refresh-lsas [lsa-type] [area-id]
 — run-manual-spf [externals-only]
```

## 5.32.2 Command Descriptions

### 5.32.2.1 Show LDP Commands



**Note:** The command outputs in this chapter are examples only; actual displays may differ depending on supported functionality and user configuration.

## bindings

|                    |                                              |
|--------------------|----------------------------------------------|
| <b>Syntax</b>      | <b>bindings</b>                              |
| <b>Context</b>     | show>router>ldp                              |
| <b>Description</b> | This command shows LDP bindings information. |
| <b>Output</b>      |                                              |

**Sample Output**

```
*A:Dut-A# show router ldp bindings
=====
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 --
10.20.1.2:0 --
--
10.20.1.1/32 262143U --
10.20.1.3:0 --
--
10.20.1.2/32 -- 262143
10.20.1.2:0 1/1/1
10.10.1.2
10.20.1.2/32 262141U 262140
10.20.1.3:0 --
--
10.20.1.3/32 262140U 262140
10.20.1.2:0 --
--
10.20.1.3/32 -- 262143
10.20.1.3:0 1/1/2
10.10.2.3
10.20.1.4/32 262139N 262139
10.20.1.2:0 1/1/1
10.10.1.2
```

|              |         |        |
|--------------|---------|--------|
| 10.20.1.4/32 | 262139U | 262139 |
| 10.20.1.3:0  | --      |        |
| --           |         |        |
| 10.20.1.5/32 | 262138U | 262137 |
| 10.20.1.2:0  | --      |        |
| --           |         |        |
| 10.20.1.5/32 | 262138N | 262137 |
| 10.20.1.3:0  | 1/1/2   |        |
| 10.10.2.3    |         |        |
| 10.20.1.6/32 | 262135N | 262135 |
| 10.20.1.2:0  | 1/1/1   |        |
| 10.10.1.2    |         |        |
| 10.20.1.6/32 | 262135U | 262135 |
| 10.20.1.3:0  | --      |        |
| --           |         |        |

-----  
No. of IPv4 Prefix Bindings: 12  
=====

=====

#### LDP IPv6 Prefix Bindings

=====

| Prefix<br>Peer<br>EgrNextHop | IngLbl<br>EgrIntf/LspId | EgrLbl |
|------------------------------|-------------------------|--------|
| -----                        | -----                   | -----  |
| 3ffe::a14:101/128            | 262142U                 | --     |
| 3ffe::a14:102[0]             | --                      |        |
| --                           |                         |        |
| 3ffe::a14:101/128            | 262142U                 | --     |
| 3ffe::a14:103[0]             | --                      |        |
| --                           |                         |        |
| 3ffe::a14:102/128            | --                      | 262142 |
| 3ffe::a14:102[0]             | 1/1/1                   |        |
| fe80::12                     |                         |        |
| 3ffe::a14:102/128            | 262136U                 | 262138 |
| 3ffe::a14:103[0]             | --                      |        |
| --                           |                         |        |
| 3ffe::a14:103/128            | 262137U                 | 262138 |
| 3ffe::a14:102[0]             | --                      |        |
| --                           |                         |        |
| 3ffe::a14:103/128            | --                      | 262142 |
| 3ffe::a14:103[0]             | 1/1/2                   |        |
| fe80::23                     |                         |        |
| 3ffe::a14:104/128            | 262132N                 | 262134 |
| 3ffe::a14:102[0]             | 1/1/1                   |        |
| fe80::12                     |                         |        |

|                   |         |        |
|-------------------|---------|--------|
| 3ffe::a14:104/128 | 262132U | 262134 |
| 3ffe::a14:103[0]  | --      |        |
| --                |         |        |
| 3ffe::a14:105/128 | 262134U | 262133 |
| 3ffe::a14:102[0]  | --      |        |
| --                |         |        |
| 3ffe::a14:105/128 | 262134N | 262132 |
| 3ffe::a14:103[0]  | 1/1/2   |        |
| fe80::23          |         |        |
| 3ffe::a14:106/128 | 262133N | 262132 |
| 3ffe::a14:102[0]  | 1/1/1   |        |
| fe80::12          |         |        |
| 3ffe::a14:106/128 | 262133U | 262133 |
| 3ffe::a14:103[0]  | --      |        |
| --                |         |        |

-----  
No. of IPv6 Prefix Bindings: 12  
=====

LDP Generic IPv4 P2MP Bindings  
=====

| P2MP-Id  | Interface   | IngLbl | EgrLbl |
|----------|-------------|--------|--------|
| RootAddr | EgrIf/LspId |        |        |
| EgrNH    |             |        |        |
| Peer     |             |        |        |

-----  
No Matching Entries Found  
=====

LDP Generic IPv6 P2MP Bindings  
=====

| P2MP-Id  | Interface   | IngLbl | EgrLbl |
|----------|-------------|--------|--------|
| RootAddr | EgrIf/LspId |        |        |
| EgrNH    |             |        |        |
| Peer     |             |        |        |

-----  
No Matching Entries Found  
=====

LDP In-Band-SSM IPv4 P2MP Bindings  
=====

| Source   | Interface   | IngLbl | EgrLbl |
|----------|-------------|--------|--------|
| Group    | EgrIf/LspId |        |        |
| RootAddr |             |        |        |
| EgrNH    |             |        |        |
| Peer     |             |        |        |

-----  
No Matching Entries Found  
=====

LDP In-Band-SSM IPv6 P2MP Bindings  
=====

Source



```

Group
RootAddr
EgrNH
Peer
Interface
EgrIf/LspId
IngLbl
EgrLbl

No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv4 P2MP Bindings
=====

Source
Group
RootAddr
EgrNH
Peer
RD
Interface
EgrIf/LspId
IngLbl
EgrLbl

No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv6 P2MP Bindings
=====

Source
Group
RootAddr
EgrNH
Peer
RD
Interface
EgrIf/LspId
IngLbl
EgrLbl

No Matching Entries Found
=====

LDP Service FEC 128 Bindings
=====

Type
Peer
VCId
SvcId
SDPID
IngLbl
LMTU
EgrLbl
RMTU

No Matching Entries Found
=====

LDP Service FEC 129 Bindings
=====

SAII
TAII
Peer
AGII
Type
SvcId
IngLbl
EgrLbl
LMTU
RMTU
SDPID

No Matching Entries Found
=====

*A:Dut-A#

*A:Dut-A# show router ldp bindings 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
=====

Prefix : 10.20.1.1/32

Peer : 10.20.1.2:0
Ing Lbl : 262143U Egr Lbl : --
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None Ing. Flags : None

Prefix : 10.20.1.1/32

Peer : 10.20.1.3:0
Ing Lbl : 262143U Egr Lbl : --
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None Ing. Flags : None

Prefix : 10.20.1.2/32

Peer : 10.20.1.2:0
Ing Lbl : -- Egr Lbl : 262143
Egr Int/LspId : 1/1/1
EgrNextHop : 10.10.1.2
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.1.1
Metric : 1000 Mtu : 1500

Prefix : 10.20.1.2/32

Peer : 10.20.1.3:0
Ing Lbl : 262141U Egr Lbl : 262140
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None Ing. Flags : None
Egr If Name : n/a

Prefix : 10.20.1.3/32

Peer : 10.20.1.2:0
Ing Lbl : 262140U Egr Lbl : 262140
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None Ing. Flags : None
Egr If Name : n/a

Prefix : 10.20.1.3/32

Peer : 10.20.1.3:0
Ing Lbl : -- Egr Lbl : 262143
Egr Int/LspId : 1/1/2
EgrNextHop : 10.10.2.3
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.2.1
Metric : 1000 Mtu : 1500

Prefix : 10.20.1.4/32
```

```

Peer : 10.20.1.2:0
Ing Lbl : 262139N Egr Lbl : 262139
Egr Int/LspId : 1/1/1
EgrNextHop : 10.10.1.2
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.1.1
Metric : 2000 Mtu : 1500

Prefix : 10.20.1.4/32

Peer : 10.20.1.3:0
Ing Lbl : 262139U Egr Lbl : 262139
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None Ing. Flags : None
Egr If Name : n/a

Prefix : 10.20.1.5/32

Peer : 10.20.1.2:0
Ing Lbl : 262138U Egr Lbl : 262137
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None Ing. Flags : None
Egr If Name : n/a

Prefix : 10.20.1.5/32

Peer : 10.20.1.3:0
Ing Lbl : 262138N Egr Lbl : 262137
Egr Int/LspId : 1/1/2
EgrNextHop : 10.10.2.3
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.2.1
Metric : 2000 Mtu : 1500

Prefix : 10.20.1.6/32

Peer : 10.20.1.2:0
Ing Lbl : 262135N Egr Lbl : 262135
Egr Int/LspId : 1/1/1
EgrNextHop : 10.10.1.2
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.1.1
Metric : 3000 Mtu : 1500

Prefix : 10.20.1.6/32

Peer : 10.20.1.3:0
Ing Lbl : 262135U Egr Lbl : 262135
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None Ing. Flags : None
Egr If Name : n/a
=====
No. of IPv4 Prefix Bindings: 12
=====
=====

```

## LDP IPv6 Prefix Bindings

```

=====
Prefix : 3ffe::a14:101/128

Peer : 3ffe::a14:102[0]
Ing Lbl : 262142U Egr Lbl : --
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None Ing. Flags : None

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::a14:102[0]
Ing Lbl : -- Egr Lbl : 262142
Egr Int/LspId : 1/1/1
EgrNextHop : fe80::12
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.1.1
Metric : 1000 Mtu : 1500

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::a14:102[0]
Ing Lbl : 262137U 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::a14:102[0]
Ing Lbl : 262132N Egr Lbl : 262134
Egr Int/LspId : 1/1/1
EgrNextHop : fe80::12
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.1.1
Metric : 2000 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::a14:102[0]
Ing Lbl : 262134U Egr Lbl : 262133
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::a14:102[0]
Ing Lbl : 262133N Egr Lbl : 262132
Egr Int/LspId : 1/1/1
EgrNextHop : fe80::12
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.1.1
Metric : 3000 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: 12
=====

```

```

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 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 --
```

|              |         |        |
|--------------|---------|--------|
| 10.20.1.2:0  | --      |        |
| --           |         |        |
| 10.20.1.1/32 | 262143U | --     |
| 10.20.1.3:0  | --      |        |
| --           |         |        |
| 10.20.1.2/32 | --      | 262143 |
| 10.20.1.2:0  | 1/1/1   |        |
| 10.10.1.2    |         |        |
| 10.20.1.2/32 | 262141U | 262140 |
| 10.20.1.3:0  | --      |        |
| --           |         |        |
| 10.20.1.3/32 | 262140U | 262140 |
| 10.20.1.2:0  | --      |        |
| --           |         |        |
| 10.20.1.3/32 | --      | 262143 |
| 10.20.1.3:0  | 1/1/2   |        |
| 10.10.2.3    |         |        |
| 10.20.1.4/32 | 262139N | 262139 |
| 10.20.1.2:0  | 1/1/1   |        |
| 10.10.1.2    |         |        |
| 10.20.1.4/32 | 262139U | 262139 |
| 10.20.1.3:0  | --      |        |
| --           |         |        |
| 10.20.1.5/32 | 262138U | 262137 |
| 10.20.1.2:0  | --      |        |
| --           |         |        |
| 10.20.1.5/32 | 262138N | 262137 |
| 10.20.1.3:0  | 1/1/2   |        |
| 10.10.2.3    |         |        |
| 10.20.1.6/32 | 262135N | 262135 |
| 10.20.1.2:0  | 1/1/1   |        |
| 10.10.1.2    |         |        |
| 10.20.1.6/32 | 262135U | 262135 |
| 10.20.1.3:0  | --      |        |
| --           |         |        |

-----  
No. of IPv4 Prefix Bindings: 12

=====

LDP Generic IPv4 P2MP Bindings

=====

| P2MP-Id  | Interface   | IngLbl | EgrLbl |
|----------|-------------|--------|--------|
| RootAddr | EgrIf/LspId |        |        |
| EgrNH    |             |        |        |
| 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 Interface IngLbl EgrLbl
EgrNH EgrIf/LspId
Peer

No Matching Entries Found
=====
*A:Dut-A#

*A:Dut-A# show router ldp bindings ipv6
=====
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 IPv6 Prefix Bindings
=====
Prefix IngLbl EgrLbl
Peer EgrIntf/LspId
EgrNextHop

3ffe::a14:101/128 262142U --
3ffe::a14:102[0] --
--

3ffe::a14:101/128 262142U --
3ffe::a14:103[0] --
--

3ffe::a14:102/128 -- 262142
3ffe::a14:102[0] 1/1/1
fe80::12

3ffe::a14:102/128 262136U 262138
3ffe::a14:103[0] --
--

```



|                   |         |        |
|-------------------|---------|--------|
| 3ffe::a14:103/128 | 262137U | 262138 |
| 3ffe::a14:102[0]  | --      |        |
| --                |         |        |
| 3ffe::a14:103/128 | --      | 262142 |
| 3ffe::a14:103[0]  | 1/1/2   |        |
| fe80::23          |         |        |
| 3ffe::a14:104/128 | 262132N | 262134 |
| 3ffe::a14:102[0]  | 1/1/1   |        |
| fe80::12          |         |        |
| 3ffe::a14:104/128 | 262132U | 262134 |
| 3ffe::a14:103[0]  | --      |        |
| --                |         |        |
| 3ffe::a14:105/128 | 262134U | 262133 |
| 3ffe::a14:102[0]  | --      |        |
| --                |         |        |
| 3ffe::a14:105/128 | 262134N | 262132 |
| 3ffe::a14:103[0]  | 1/1/2   |        |
| fe80::23          |         |        |
| 3ffe::a14:106/128 | 262133N | 262132 |
| 3ffe::a14:102[0]  | 1/1/1   |        |
| fe80::12          |         |        |
| 3ffe::a14:106/128 | 262133U | 262133 |
| 3ffe::a14:103[0]  | --      |        |
| --                |         |        |

-----  
No. of IPv6 Prefix Bindings: 12

=====

LDP Generic IPv6 P2MP Bindings

=====

| P2MP-Id  | Interface   | IngLbl | EgrLbl |
|----------|-------------|--------|--------|
| RootAddr | EgrIf/LspId |        |        |
| EgrNH    |             |        |        |
| Peer     |             |        |        |

-----  
No Matching Entries Found

=====

LDP In-Band-SSM IPv6 P2MP Bindings

=====

| Source   | Interface   | IngLbl | EgrLbl |
|----------|-------------|--------|--------|
| Group    | EgrIf/LspId |        |        |
| RootAddr |             |        |        |
| EgrNH    |             |        |        |
| Peer     |             |        |        |

-----  
No Matching Entries Found

=====

LDP In-Band-VPN-SSM IPv6 P2MP Bindings

=====

```

Source
Group
RootAddr
EgrNH
Peer
RD
Interface
EgrIf/LspId
IngLbl
EgrLbl

No Matching Entries Found
=====
*A:Dut-A#

*A:Dut-A# show router ldp bindings label-type start-label 262100 end-label 262300
egress-label
=====
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
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP IPv4 Prefix Bindings
=====
Prefix
Peer
EgrNextHop
IngLbl
EgrIntf/LspId
EgrLbl

10.20.1.2/32
10.20.1.2:0
10.10.1.2
--
1/1/1
262143

10.20.1.2/32
10.20.1.3:0
--
262141U
--
262140

10.20.1.3/32
10.20.1.2:0
--
262140U
--
262140

10.20.1.3/32
10.20.1.3:0
10.10.2.3
--
1/1/2
262143

10.20.1.4/32
10.20.1.2:0
10.10.1.2
262139N
1/1/1
262139

10.20.1.4/32
10.20.1.3:0
--
262139U
--
262139

10.20.1.5/32
10.20.1.2:0
--
262138U
--
262137

10.20.1.5/32
10.20.1.3:0
10.10.2.3
262138N
1/1/2
262137

10.20.1.6/32
10.20.1.2:0
262135N
1/1/1
262135

```

```

10.10.1.2

10.20.1.6/32 262135U 262135
10.20.1.3:0 --
--

No. of IPv4 Prefix Bindings: 10
=====
LDP IPv6 Prefix Bindings
=====
Prefix IngLbl EgrLbl
Peer EgrIntf/LspId
EgrNextHop

3ffe::a14:102/128 -- 262142
3ffe::a14:102[0] 1/1/1
fe80::12

3ffe::a14:102/128 262136U 262138
3ffe::a14:103[0] --
--

3ffe::a14:103/128 262137U 262138
3ffe::a14:102[0] --
--

3ffe::a14:103/128 -- 262142
3ffe::a14:103[0] 1/1/2
fe80::23

3ffe::a14:104/128 262132N 262134
3ffe::a14:102[0] 1/1/1
fe80::12

3ffe::a14:104/128 262132U 262134
3ffe::a14:103[0] --
--

3ffe::a14:105/128 262134U 262133
3ffe::a14:102[0] --
--

3ffe::a14:105/128 262134N 262132
3ffe::a14:103[0] 1/1/2
fe80::23

3ffe::a14:106/128 262133N 262132
3ffe::a14:102[0] 1/1/1
fe80::12

3ffe::a14:106/128 262133U 262133
3ffe::a14:103[0] --
--

No. of IPv6 Prefix Bindings: 10
=====

```

```

=====
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:Dut-A#
*A:Dut-A#
*A:Dut-A# show router ldp bindings label-type start-label 262100 end-label 262300
egress-label ipv6
=====
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
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP IPv6 Prefix Bindings
=====
Prefix IngLbl EgrLbl
Peer EgrIntf/LspId
EgrNextHop

3ffe::a14:102/128 -- 262142
3ffe::a14:102[0] 1/1/1
fe80::12

3ffe::a14:102/128 262136U 262138
3ffe::a14:103[0] --
--

3ffe::a14:103/128 262137U 262138
3ffe::a14:102[0] --
--

3ffe::a14:103/128 -- 262142
3ffe::a14:103[0] 1/1/2
fe80::23

3ffe::a14:104/128 262132N 262134
3ffe::a14:102[0] 1/1/1
fe80::12

3ffe::a14:104/128 262132U 262134
3ffe::a14:103[0] --
--

3ffe::a14:105/128 262134U 262133

```

```
3ffe::a14:102[0] --
--

3ffe::a14:105/128 262134N 262132
3ffe::a14:103[0] 1/1/2
fe80::23

3ffe::a14:106/128 262133N 262132
3ffe::a14:102[0] 1/1/1
fe80::12

3ffe::a14:106/128 262133U 262133
3ffe::a14:103[0] --
--
```

-----  
No. of IPv6 Prefix Bindings: 10  
=====

\*A:Dut-A#

\*A:Dut-A# show router ldp bindings prefixes prefix 3ffe::a14:104/128

=====

```
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
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
```

=====

```
LDP IPv6 Prefix Bindings
```

=====

| Prefix            | IngLbl        | EgrLbl |
|-------------------|---------------|--------|
| Peer              | EgrIntf/LspId |        |
| EgrNextHop        |               |        |
| 3ffe::a14:104/128 | 262132N       | 262134 |
| 3ffe::a14:102[0]  | 1/1/1         |        |
| fe80::12          |               |        |
| 3ffe::a14:104/128 | 262132U       | 262134 |
| 3ffe::a14:103[0]  | --            |        |
| --                |               |        |

-----

No. of IPv6 Prefix Bindings: 2

=====

```
*A:Dut-A# show router ldp bindings prefixes prefix 3ffe::a14:104/128 summary
```

```
No. of IPv6 Prefix Bindings: 2
```

```
*A:Dut-A# show router ldp bindings prefixes prefix 3ffe::a14:104/128 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
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
```

=====

```
LDP IPv6 Prefix Bindings
```

-----

```
Prefix : 3ffe::a14:104/128
```

```

Peer : 3ffe::a14:102[0]
Ing Lbl : 262132N Egr Lbl : 262134
Egr Int/LspId : 1/1/1
EgrNextHop : fe80::12
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.1.1
Metric : 2000 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
=====
No. of IPv6 Prefix Bindings: 2
=====
*A:Dut-A#

*A:Dut-A# show router ldp bindings prefixes prefix 3ffe::a14:104/128 session
3ffe::a14:103
=====
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
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP IPv6 Prefix Bindings
=====
Prefix IngLbl EgrLbl
Peer EgrIntf/LspId
EgrNextHop

3ffe::a14:104/128 262132U 262134
3ffe::a14:103[0] --
--

No. of IPv6 Prefix Bindings: 1
=====
*A:Dut-A# show router ldp bindings prefixes prefix 3ffe::a14:104/128 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
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP IPv6 Prefix Bindings
=====
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
=====
No. of IPv6 Prefix Bindings: 1
=====
*A:Dut-A# show router ldp bindings prefixes prefix 3ffe::a14:104/128 session
3ffe::a14:103 summary
 No. of IPv6 Prefix Bindings: 1
*A:Dut-A#

*A:Dut-A# show router ldp bindings session 3ffe::a14:103 ipv6
=====
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 IPv6 Prefix Bindings
=====
Prefix IngLbl EgrLbl
Peer EgrIntf/LspId
EgrNextHop

3ffe::a14:101/128 262142U --
3ffe::a14:103[0] --
--
3ffe::a14:102/128 262136U 262138
3ffe::a14:103[0] --
--
3ffe::a14:103/128 -- 262142
3ffe::a14:103[0] 1/1/2
fe80::23
3ffe::a14:104/128 262132U 262134
3ffe::a14:103[0] --
--
3ffe::a14:105/128 262134N 262132
3ffe::a14:103[0] 1/1/2
fe80::23
3ffe::a14:106/128 262133U 262133
3ffe::a14:103[0] --
--

No. of IPv6 Prefix Bindings: 6
=====
=====

```

```

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
RootAddr Interface IngLbl EgrLbl
EgrNH EgrIf/LspId
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
=====
*A:Dut-A#

*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<br>Peer<br>EgrNextHop | IngLbl<br>EgrIntf/LspId | EgrLbl |
|------------------------------|-------------------------|--------|
| 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<br>RootAddr<br>EgrNH<br>Peer | Interface<br>EgrIf/LspId | IngLbl | EgrLbl |
|--------------------------------------|--------------------------|--------|--------|
| No Matching Entries Found            |                          |        |        |

=====

LDP In-Band-SSM IPv4 P2MP Bindings

=====

| Source<br>Group<br>RootAddr<br>EgrNH<br>Peer | Interface<br>EgrIf/LspId | IngLbl | EgrLbl |
|----------------------------------------------|--------------------------|--------|--------|
| No Matching Entries Found                    |                          |        |        |

=====

LDP In-Band-VPN-SSM IPv4 P2MP Bindings

=====

| Source<br>Group | RD |
|-----------------|----|
|                 |    |

```

RootAddr Interface IngLbl EgrLbl
EgrNH EgrIf/LspId
Peer

No Matching Entries Found
=====
*A:Dut-A#

*A:Dut-C# show router ldp bindings active prefixes prefix 10.20.1.5/32
=====
LDP Bindings (IPv4 LSR ID 10.20.1.3)
 (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
 (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
 (C) - FEC resolved with class-based-forwarding
=====
LDP IPv4 Prefix Bindings (Active)
=====
Prefix Op IngLbl EgrLbl
EgrNextHop EgrIf/LspId

10.20.1.5/32 Push -- 262135
10.10.5.5 2/1/1
10.20.1.5/32 Push -- BU
10.20.1.5 LspId 524300
10.20.1.5/32 Swap 262126 262135
10.10.5.5 2/1/1
10.20.1.5/32 Swap 262126 BU
10.20.1.5 LspId 524300

No. of IPv4 Prefix Active Bindings: 4
=====

*A:Dut-C# show router ldp bindings active p2mp detail
=====
LDP Bindings (IPv4 LSR ID 10.20.1.3)
 (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, e - Label ELC
=====
LDP Generic IPv4 P2MP Bindings (Active)
=====
P2MP Type : 1 P2MP-Id : 1
Root-Addr : 10.20.1.3

Op : Push
Ing Lbl : --

```

```

Egr Lbl : 262139
Egr Int/LspId : 1/1/2:0
EgrNextHop : 10.180.3.2
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.3.3
Metric : 1 Mtu : 1496

P2MP Type : 1 P2MP-Id : 1
Root-Addr : 10.20.1.3

Op : Push
Ing Lbl : --
Egr Lbl : 262137
Egr Int/LspId : 2/1/2:0
EgrNextHop : 10.180.11.4
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.11.3
Metric : 1 Mtu : 1496

P2MP Type : 1 P2MP-Id : 2
Root-Addr : 10.20.1.3

Op : Push
Ing Lbl : --
Egr Lbl : 262138
Egr Int/LspId : 1/1/2:0
EgrNextHop : 10.180.3.2
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.3.3
Metric : 1 Mtu : 1496

P2MP Type : 1 P2MP-Id : 2
Root-Addr : 10.20.1.3

Op : Push
Ing Lbl : --
Egr Lbl : 262136
Egr Int/LspId : 2/1/2:0
EgrNextHop : 10.180.11.4
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.11.3
Metric : 1 Mtu : 1496

P2MP Type : 1 P2MP-Id : 3
Root-Addr : 10.20.1.3

*A:Dut-C#
*A:Dut-C# show router ldp bindings p2mp detail
=====
LDP Bindings (IPv4 LSR ID 10.20.1.3)
 (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, e - Label ELC
=====
LDP Generic IPv4 P2MP Bindings
=====

```

---

```
P2MP Type : 1 P2MP-Id : 1
Root-Addr : 10.20.1.3

Peer : 10.20.1.2:0
Ing Lbl : --
Egr Lbl : 262139
Egr Int/LspId : 1/1/2:0
EgrNextHop : 10.180.3.2
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.3.3
Metric : 1 Mtu : 1496

P2MP Type : 1 P2MP-Id : 1
Root-Addr : 10.20.1.3

Peer : 10.20.1.4:0
Ing Lbl : --
Egr Lbl : 262137
Egr Int/LspId : 2/1/2:0
EgrNextHop : 10.180.11.4
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.11.3
Metric : 1 Mtu : 1496

P2MP Type : 1 P2MP-Id : 2
Root-Addr : 10.20.1.3

Peer : 10.20.1.2:0
Ing Lbl : --
Egr Lbl : 262138
Egr Int/LspId : 1/1/2:0
EgrNextHop : 10.180.3.2
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.3.3
Metric : 1 Mtu : 1496

P2MP Type : 1 P2MP-Id : 2
Root-Addr : 10.20.1.3

Peer : 10.20.1.4:0
Ing Lbl : --
Egr Lbl : 262136
Egr Int/LspId : 2/1/2:0
EgrNextHop : 10.180.11.4
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.11.3
Metric : 1 Mtu : 1496

P2MP Type : 1 P2MP-Id : 3
Root-Addr : 10.20.1.3

Peer : 10.20.1.2:0
Ing Lbl : --
Egr Lbl : 262137
Egr Int/LspId : 1/1/2:0
EgrNextHop : 10.180.3.2
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.3.3
Metric : 1 Mtu : 1496
```

```

P2MP Type : 1 P2MP-Id : 3
Root-Addr : 10.20.1.3

Peer : 10.20.1.4:0
Ing Lbl : --
Egr Lbl : 262135
Egr Int/LspId : 2/1/2:0
EgrNextHop : 10.180.11.4
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.11.3
Metric : 1 Mtu : 1496

P2MP Type : 1 P2MP-Id : 4
Root-Addr : 10.20.1.3

Peer : 10.20.1.2:0
Ing Lbl : --
Egr Lbl : 262136
Egr Int/LspId : 1/1/2:0
EgrNextHop : 10.180.3.2
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.3.3
Metric : 1 Mtu : 1496

P2MP Type : 1 P2MP-Id : 4
Root-Addr : 10.20.1.3

Peer : 10.20.1.4:0
Ing Lbl : --
Egr Lbl : 262134
Egr Int/LspId : 2/1/2:0
EgrNextHop : 10.180.11.4
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.11.3
Metric : 1 Mtu : 1496

P2MP Type : 1 P2MP-Id : 5
Root-Addr : 10.20.1.3

Peer : 10.20.1.2:0
Ing Lbl : --
Egr Lbl : 262135
Egr Int/LspId : 1/1/2:0
EgrNextHop : 10.180.3.2
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.3.3
Metric : 1 Mtu : 1496

P2MP Type : 1 P2MP-Id : 5
Root-Addr : 10.20.1.3

Peer : 10.20.1.4:0
Ing Lbl : --
Egr Lbl : 262133
Egr Int/LspId : 2/1/2:0
EgrNextHop : 10.180.11.4
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.180.11.3

```

```

Metric : 1 Mtu : 1496

P2MP Type : 1 P2MP-Id : 6
Root-Addr : 10.20.1.3

Peer : 10.20.1.2:0
Ing Lbl : --
Egr Lbl : 262134
Egr Int/LspId : 1/1/2:0
EgrNextHop : 10.180.3.2
Egr. Flags : None Ing. Flags : None

```

## active

**Syntax**

**active egress-lsp** [*tunnel-id*]  
**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 p2mp-id** *identifier* *root ip-address* [**rd** *rd*] [**summary** | **detail**] [**egress-if** *port-id*]  
*inner-root ip-address*  
**active p2mp p2mp-id** *identifier* *root ip-address* [**rd** *rd*] [**summary** | **detail**] [**egress-lsp** *tunnel-id*] *inner-root ip-address*  
**active p2mp p2mp-id** *identifier* *root ip-address* [**rd** *rd*] [**summary** | **detail**] [**egress-nh** *ip-address*] *inner-root 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* **root** *ip-address* [**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 display 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.  
**opaque-type** *opaque-type* — Specifies the type of a Multi-Point Opaque Value Element.  
**Values** generic, ssm, vpn-ssm, recursive-ssm  
**egress-lsp** *tunnel-id* — Displays LDP active bindings by matching on the egress RSVP-TE LSP *tunnel-id* for LDP FECs which are tunneled over a 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.  
**Values**

|              |                                       |
|--------------|---------------------------------------|
| ipv4-address | - a.b.c.d                             |
| ipv6-address | - x:x:x:x:x:x:x (eight 16-bit pieces) |
|              | x:x:x:x:x:x.d.d.d.d                   |
|              | x - [0 to FFFF]H                      |
|              | d - [0 to 255]D                       |

**egress-if** *port-id* — Displays LDP active bindings by matching egress-if.  
**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.

**Values**

ipv4-address      - a.b.c.d  
 ipv6-address      - x:x:x:x:x:x:x (eight 16-bit pieces)  
                       x:x:x:x:x:x:d.d.d.d  
                       x - [0 to FFFF]H  
                       d - [0 to 255]D

**p2mp-id *identifier*** — Displays LDP active P2MP identifier bindings.

**Values**      0 to 4294 967 295

**group *mcast-address*** — Displays the P2MP group multicast address bindings.

**root *ip-address*** — Displays root IP address information.

**rd *rd*** — Displays information for the route distinguisher.

**Values**      *ip-addr:comm-val* | *2byte-asnumber:ext-comm-val* | *4byte-asnumber:comm-val*

**prefix *ip-prefix/ip-prefix-length*** — Specifies information for the specified IP prefix and mask length.

**Values**

ipv4-address      - a.b.c.d  
 ipv6-address      - x:x:x:x:x:x:x (eight 16-bit pieces)  
                       x:x:x:x:x:x:d.d.d.d  
                       x - [0 to FFFF]H  
                       d - [0 to 255]D

## Output

### Sample Output

```
*A:Dut-C# 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
 (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<br>EgrNextHop       | Op<br>EgrIf/LspId | IngLbl | EgrLbl |
|----------------------------|-------------------|--------|--------|
| 10.20.1.1/32<br>10.10.2.1  | Push<br>1/1/1     | --     | 262143 |
| 10.20.1.1/32<br>10.10.2.1  | Swap<br>1/1/1     | 262141 | 262143 |
| 10.20.1.2/32<br>10.10.12.2 | Push<br>lag-1     | --     | 262143 |
| 10.20.1.2/32<br>10.10.12.2 | Swap<br>lag-1     | 262140 | 262143 |
| 10.20.1.3/32<br>--         | Pop<br>--         | 262143 | --     |
| 10.20.1.4/32<br>10.10.11.4 | Push<br>2/1/2     | --     | 262143 |
| 10.20.1.4/32<br>10.10.11.4 | Swap<br>2/1/2     | 262139 | 262143 |
| 10.20.1.5/32<br>10.10.5.5  | Push<br>2/1/1     | --     | 262143 |
| 10.20.1.5/32<br>10.10.5.5  | Swap<br>2/1/1     | 262137 | 262143 |
| 10.20.1.6/32<br>10.10.11.4 | Push<br>2/1/2     | --     | 262137 |
| 10.20.1.6/32<br>10.10.11.4 | Swap<br>2/1/2     | 262135 | 262137 |

-----  
No. of IPv4 Prefix Active Bindings: 11  
=====

=====  
LDP IPv6 Prefix Bindings (Active)  
=====

| Prefix<br>EgrNextHop           | Op<br>EgrIf/LspId | IngLbl | EgrLbl |
|--------------------------------|-------------------|--------|--------|
| 3ffe::a14:101/128<br>fe80::21  | Push<br>1/1/1     | --     | 262142 |
| 3ffe::a14:101/128<br>fe80::21  | Swap<br>1/1/1     | 262136 | 262142 |
| 3ffe::a14:102/128<br>fe80::122 | Push<br>lag-1     | --     | 262142 |
| 3ffe::a14:102/128<br>fe80::122 | Swap<br>lag-1     | 262138 | 262142 |
| 3ffe::a14:103/128              | Pop               | 262142 | --     |

```
--
--
3ffe::a14:104/128 Push -- 262142
fe80::114 2/1/2

3ffe::a14:104/128 Swap 262134 262142
fe80::114 2/1/2

3ffe::a14:105/128 Push -- 262142
fe80::55 2/1/1

3ffe::a14:105/128 Swap 262132 262142
fe80::55 2/1/1

3ffe::a14:106/128 Push -- 262136
fe80::114 2/1/2

3ffe::a14:106/128 Swap 262133 262136
fe80::114 2/1/2

No. of IPv6 Prefix Active Bindings: 11
=====

=====
LDP Generic IPv4 P2MP Bindings (Active)
=====
P2MP-Id Interface
RootAddr Op IngLbl EgrLbl
EgrNH EgrIf/LspId

No Matching Entries Found
=====

=====
LDP Generic IPv6 P2MP Bindings (Active)
=====
P2MP-Id Interface
RootAddr Op IngLbl EgrLbl
EgrNH EgrIf/LspId

No Matching Entries Found
=====

=====
LDP In-Band-SSM IPv4 P2MP Bindings (Active)
=====
Source
Group Interface
RootAddr Op IngLbl EgrLbl
EgrNH EgrIf/LspId

No Matching Entries Found
=====

=====
LDP In-Band-SSM IPv6 P2MP Bindings (Active)
=====
Source
```

```
Group Interface
RootAddr Op IngLbl EgrLbl
EgrNH EgrIf/LspId

No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv4 P2MP Bindings (Active)
=====
Source
Group RD Op
RootAddr Interface IngLbl EgrLbl
EgrNH EgrIf/LspId

No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv6 P2MP Bindings (Active)
=====
Source
Group RD Op
RootAddr Interface IngLbl EgrLbl
EgrNH EgrIf/LspId

No Matching Entries Found
=====

*A:Dut-C# 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

Prefix : 10.20.1.2/32
Op : Push
Ing Lbl : -- Egr Lbl : 262143
Egr Int/LspId : lag-1
EgrNextHop : 10.10.12.2
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.12.3
Metric : 333 Mtu : 1500

Prefix : 10.20.1.2/32
Op : Swap
Ing Lbl : 262140 Egr Lbl : 262143
Egr Int/LspId : lag-1
EgrNextHop : 10.10.12.2
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.12.3
Metric : 333 Mtu : 1500

Prefix : 10.20.1.3/32
Op : Pop
Ing Lbl : 262143 Egr Lbl : --
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None Ing. Flags : None

Prefix : 10.20.1.4/32
Op : Push
Ing Lbl : -- Egr Lbl : 262143
Egr Int/LspId : 2/1/2
EgrNextHop : 10.10.11.4
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.11.3
Metric : 1000 Mtu : 1500

Prefix : 10.20.1.4/32
Op : Swap
Ing Lbl : 262139 Egr Lbl : 262143
Egr Int/LspId : 2/1/2
EgrNextHop : 10.10.11.4
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.11.3
Metric : 1000 Mtu : 1500

Prefix : 10.20.1.5/32
Op : Push
Ing Lbl : -- Egr Lbl : 262143
Egr Int/LspId : 2/1/1
EgrNextHop : 10.10.5.5
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.5.3
Metric : 1000 Mtu : 1500

Prefix : 10.20.1.5/32
Op : Swap
Ing Lbl : 262137 Egr Lbl : 262143
Egr Int/LspId : 2/1/1
EgrNextHop : 10.10.5.5

```

```

Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.5.3
Metric : 1000 Mtu : 1500

Prefix : 10.20.1.6/32
Op : Push
Ing Lbl : -- Egr Lbl : 262137
Egr Int/LspId : 2/1/2
EgrNextHop : 10.10.11.4
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.11.3
Metric : 2000 Mtu : 1500

Prefix : 10.20.1.6/32
Op : Swap
Ing Lbl : 262135 Egr Lbl : 262137
Egr Int/LspId : 2/1/2
EgrNextHop : 10.10.11.4
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.11.3
Metric : 2000 Mtu : 1500
=====
No. of IPv4 Prefix Active Bindings: 11
=====

=====
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

Prefix : 3ffe::a14:102/128
Op : Push
Ing Lbl : -- Egr Lbl : 262142
Egr Int/LspId : lag-1
EgrNextHop : fe80::122
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.12.3
Metric : 333 Mtu : 1500

Prefix : 3ffe::a14:102/128
Op : Swap
Ing Lbl : 262138 Egr Lbl : 262142

```

```

Egr Int/LspId : lag-1
EgrNextHop : fe80::122
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.12.3
Metric : 333 Mtu : 1500

Prefix : 3ffe::a14:103/128
Op : Pop
Ing Lbl : 262142 Egr Lbl : --
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None Ing. Flags : None

Prefix : 3ffe::a14:104/128
Op : Push
Ing Lbl : -- Egr Lbl : 262142
Egr Int/LspId : 2/1/2
EgrNextHop : fe80::114
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.11.3
Metric : 1000 Mtu : 1500

Prefix : 3ffe::a14:104/128
Op : Swap
Ing Lbl : 262134 Egr Lbl : 262142
Egr Int/LspId : 2/1/2
EgrNextHop : fe80::114
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.11.3
Metric : 1000 Mtu : 1500

Prefix : 3ffe::a14:105/128
Op : Push
Ing Lbl : -- Egr Lbl : 262142
Egr Int/LspId : 2/1/1
EgrNextHop : fe80::55
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.5.3
Metric : 1000 Mtu : 1500

Prefix : 3ffe::a14:105/128
Op : Swap
Ing Lbl : 262132 Egr Lbl : 262142
Egr Int/LspId : 2/1/1
EgrNextHop : fe80::55
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.5.3
Metric : 1000 Mtu : 1500

Prefix : 3ffe::a14:106/128
Op : Push
Ing Lbl : -- Egr Lbl : 262136
Egr Int/LspId : 2/1/2
EgrNextHop : fe80::114
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.11.3
Metric : 2000 Mtu : 1500

Prefix : 3ffe::a14:106/128

```



```
Op : Swap
Ing Lbl : 262133 Egr Lbl : 262136
Egr Int/LspId : 2/1/2
EgrNextHop : fe80::114
Egr. Flags : None Ing. Flags : None
Egr If Name : ip-10.10.11.3
Metric : 2000 Mtu : 1500
=====
No. of IPv6 Prefix Active Bindings: 11
=====
LDP Generic IPv4 P2MP Bindings (Active)
=====
No Matching Entries Found
=====
LDP Generic IPv6 P2MP Bindings (Active)
=====
No Matching Entries Found
=====
LDP In-Band-SSM IPv4 P2MP Bindings (Active)
=====
No Matching Entries Found
=====
LDP In-Band-SSM IPv6 P2MP Bindings (Active)
=====
No Matching Entries Found
=====
LDP In-Band-VPN-SSM IPv4 P2MP Bindings (Active)
=====
No Matching Entries Found
=====
LDP In-Band-VPN-SSM IPv6 P2MP Bindings (Active)
=====
No Matching Entries Found
=====
*A:Dut-C# show router ldp bindings active egress-if 2/1/2
=====
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 Op IngLbl EgrLbl
EgrNextHop EgrIf/LspId

10.20.1.4/32 Push -- 262143
10.10.11.4 2/1/2
10.20.1.4/32 Swap 262139 262143
```

```

10.10.11.4 2/1/2

10.20.1.6/32 Push -- 262137
10.10.11.4 2/1/2

10.20.1.6/32 Swap 262135 262137
10.10.11.4 2/1/2

```

```

No. of IPv4 Prefix Active Bindings: 4
=====

```

```

LDP IPv6 Prefix Bindings (Active)
=====

```

| Prefix<br>EgrNextHop           | Op<br>EgrIf/LspId | IngLbl | EgrLbl |
|--------------------------------|-------------------|--------|--------|
| 3ffe::a14:104/128<br>fe80::114 | Push<br>2/1/2     | --     | 262142 |
| 3ffe::a14:104/128<br>fe80::114 | Swap<br>2/1/2     | 262134 | 262142 |
| 3ffe::a14:106/128<br>fe80::114 | Push<br>2/1/2     | --     | 262136 |
| 3ffe::a14:106/128<br>fe80::114 | Swap<br>2/1/2     | 262133 | 262136 |

```

No. of IPv6 Prefix Active Bindings: 4
=====

```

```

LDP Generic IPv4 P2MP Bindings (Active)
=====

```

| P2MP-Id<br>RootAddr<br>EgrNH | Interface<br>Op<br>EgrIf/LspId | IngLbl | EgrLbl |
|------------------------------|--------------------------------|--------|--------|
| No Matching Entries Found    |                                |        |        |

```

```

```

LDP Generic IPv6 P2MP Bindings (Active)
=====

```

| P2MP-Id<br>RootAddr<br>EgrNH | Interface<br>Op<br>EgrIf/LspId | IngLbl | EgrLbl |
|------------------------------|--------------------------------|--------|--------|
| No Matching Entries Found    |                                |        |        |

```

```

```

LDP In-Band-SSM IPv4 P2MP Bindings (Active)
=====

```

| Source<br>Group<br>RootAddr<br>EgrNH | Interface<br>Op<br>EgrIf/LspId | IngLbl | EgrLbl |
|--------------------------------------|--------------------------------|--------|--------|
| No Matching Entries Found            |                                |        |        |

```

```

```

=====
LDP In-Band-SSM IPv6 P2MP Bindings (Active)
=====
Source
Group Interface
RootAddr Op IngLbl EgrLbl
EgrNH EgrIf/LspId

No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv4 P2MP Bindings (Active)
=====
Source
Group RD Op
RootAddr Interface IngLbl EgrLbl
EgrNH EgrIf/LspId

No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv6 P2MP Bindings (Active)
=====
Source
Group RD Op
RootAddr Interface IngLbl EgrLbl
EgrNH EgrIf/LspId

No Matching Entries Found
=====

*A:Dut-C# show router ldp bindings active egress-nh 10.10.11.4
=====
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 Op IngLbl EgrLbl
EgrNextHop EgrIf/LspId

10.20.1.4/32 Push -- 262143
10.10.11.4 2/1/2

10.20.1.4/32 Swap 262139 262143
10.10.11.4 2/1/2

10.20.1.6/32 Push -- 262137
10.10.11.4 2/1/2

10.20.1.6/32 Swap 262135 262137
10.10.11.4 2/1/2

```

```

No. of IPv4 Prefix Active Bindings: 4
=====
```

```
LDP IPv6 Prefix Bindings (Active)
=====
```

| Prefix     | Op          | IngLbl | EgrLbl |
|------------|-------------|--------|--------|
| EgrNextHop | EgrIf/LspId |        |        |

```

No Matching Entries Found
=====
```

```
LDP Generic IPv4 P2MP Bindings (Active)
=====
```

| P2MP-Id  | Interface   |        |        |
|----------|-------------|--------|--------|
| RootAddr | Op          | IngLbl | EgrLbl |
| EgrNH    | EgrIf/LspId |        |        |

```

No Matching Entries Found
=====
```

```
LDP Generic IPv6 P2MP Bindings (Active)
=====
```

| P2MP-Id  | Interface   |        |        |
|----------|-------------|--------|--------|
| RootAddr | Op          | IngLbl | EgrLbl |
| EgrNH    | EgrIf/LspId |        |        |

```

No Matching Entries Found
=====
```

```
LDP In-Band-SSM IPv4 P2MP Bindings (Active)
=====
```

| Source   |             |        |        |
|----------|-------------|--------|--------|
| Group    | Interface   |        |        |
| RootAddr | Op          | IngLbl | EgrLbl |
| EgrNH    | EgrIf/LspId |        |        |

```

No Matching Entries Found
=====
```

```
LDP In-Band-SSM IPv6 P2MP Bindings (Active)
=====
```

| Source   |             |        |        |
|----------|-------------|--------|--------|
| Group    | Interface   |        |        |
| RootAddr | Op          | IngLbl | EgrLbl |
| EgrNH    | EgrIf/LspId |        |        |

```

No Matching Entries Found
=====
```

```
LDP In-Band-VPN-SSM IPv4 P2MP Bindings (Active)
=====
```

| Source   |             |        |        |
|----------|-------------|--------|--------|
| Group    | RD          | Op     |        |
| RootAddr | Interface   | IngLbl | EgrLbl |
| EgrNH    | EgrIf/LspId |        |        |

```

No Matching Entries Found
=====
```

```

=====
LDP In-Band-VPN-SSM IPv6 P2MP Bindings (Active)
=====
Source
Group RD Op
RootAddr Interface IngLbl EgrLbl
EgrNH EgrIf/LspId

No Matching Entries Found
=====

*A:Dut-C# show router ldp bindings active ipv4
=====
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 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

10.20.1.2/32 Push -- 262143
10.10.12.2 lag-1

10.20.1.2/32 Swap 262140 262143
10.10.12.2 lag-1

10.20.1.3/32 Pop 262143 --
-- --

10.20.1.4/32 Push -- 262143
10.10.11.4 2/1/2

10.20.1.4/32 Swap 262139 262143
10.10.11.4 2/1/2

10.20.1.5/32 Push -- 262143
10.10.5.5 2/1/1

10.20.1.5/32 Swap 262137 262143
10.10.5.5 2/1/1

10.20.1.6/32 Push -- 262137
10.10.11.4 2/1/2

10.20.1.6/32 Swap 262135 262137
10.10.11.4 2/1/2

```

```

No. of IPv4 Prefix Active Bindings: 11
=====
```

```
LDP Generic IPv4 P2MP Bindings (Active)
=====
```

| P2MP-Id  | Interface   |        |        |
|----------|-------------|--------|--------|
| RootAddr | Op          | IngLbl | EgrLbl |
| EgrNH    | EgrIf/LspId |        |        |

```

No Matching Entries Found
=====
```

```
LDP In-Band-SSM IPv4 P2MP Bindings (Active)
=====
```

| Source   | Interface   |        |        |
|----------|-------------|--------|--------|
| Group    | Op          | IngLbl | EgrLbl |
| RootAddr | EgrIf/LspId |        |        |
| EgrNH    |             |        |        |

```

No Matching Entries Found
=====
```

```
LDP In-Band-VPN-SSM IPv4 P2MP Bindings (Active)
=====
```

| Source   | RD          | Op     |        |
|----------|-------------|--------|--------|
| Group    | Interface   | IngLbl | EgrLbl |
| RootAddr | EgrIf/LspId |        |        |
| EgrNH    |             |        |        |

```

No Matching Entries Found
=====
```

```
*A:Dut-C# show router ldp bindings active ipv6
=====
```

```
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 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       |        |        |
| 3ffe::a14:102/128 | Push        | --     | 262142 |
| fe80::122         | lag-1       |        |        |
| 3ffe::a14:102/128 | Swap        | 262138 | 262142 |
| fe80::122         | lag-1       |        |        |

|                   |       |        |        |
|-------------------|-------|--------|--------|
| 3ffe::a14:103/128 | Pop   | 262142 | --     |
| --                | --    |        |        |
| 3ffe::a14:104/128 | Push  | --     | 262142 |
| fe80::114         | 2/1/2 |        |        |
| 3ffe::a14:104/128 | Swap  | 262134 | 262142 |
| fe80::114         | 2/1/2 |        |        |
| 3ffe::a14:105/128 | Push  | --     | 262142 |
| fe80::55          | 2/1/1 |        |        |
| 3ffe::a14:105/128 | Swap  | 262132 | 262142 |
| fe80::55          | 2/1/1 |        |        |
| 3ffe::a14:106/128 | Push  | --     | 262136 |
| fe80::114         | 2/1/2 |        |        |
| 3ffe::a14:106/128 | Swap  | 262133 | 262136 |
| fe80::114         | 2/1/2 |        |        |

-----  
No. of IPv6 Prefix Active Bindings: 11  
=====

LDP Generic IPv6 P2MP Bindings (Active)  
=====

| P2MP-Id  | Interface   |        |        |
|----------|-------------|--------|--------|
| RootAddr | Op          | IngLbl | EgrLbl |
| EgrNH    | EgrIf/LspId |        |        |

-----  
No Matching Entries Found  
=====

LDP In-Band-SSM IPv6 P2MP Bindings (Active)  
=====

| Source   |             |        |        |
|----------|-------------|--------|--------|
| Group    | Interface   |        |        |
| RootAddr | Op          | IngLbl | EgrLbl |
| EgrNH    | EgrIf/LspId |        |        |

-----  
No Matching Entries Found  
=====

LDP In-Band-VPN-SSM IPv6 P2MP Bindings (Active)  
=====

| Source   |             |        |        |
|----------|-------------|--------|--------|
| Group    | RD          | Op     |        |
| RootAddr | Interface   | IngLbl | EgrLbl |
| EgrNH    | EgrIf/LspId |        |        |

-----  
No Matching Entries Found  
=====

\*A:Dut-C# show router ldp bindings active summary

No. of IPv4 Prefix Active Bindings: 11  
No. of IPv6 Prefix Active Bindings: 11  
No. of Generic IPv4 P2MP Active Bindings: 0  
No. of Generic IPv6 P2MP Active Bindings: 0

```

No. of In-Band-SSM IPv4 P2MP Active Bindings: 0
No. of In-Band-SSM IPv6 P2MP Active Bindings: 0
No. of In-Band-VPN-SSM IPv4 P2MP Active Bindings: 0
No. of In-Band-VPN-SSM IPv6 P2MP Active Bindings: 0

```

```
*A:Dut-C#
```

```
*A:Dut-C# show router ldp bindings active prefixes
```

```

=====
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 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
10.20.1.2/32 Push -- 262143
10.10.12.2 lag-1
10.20.1.2/32 Swap 262140 262143
10.10.12.2 lag-1
10.20.1.3/32 Pop 262143 --
-- --
10.20.1.4/32 Push -- 262143
10.10.11.4 2/1/2
10.20.1.4/32 Swap 262139 262143
10.10.11.4 2/1/2
10.20.1.5/32 Push -- 262143
10.10.5.5 2/1/1
10.20.1.5/32 Swap 262137 262143
10.10.5.5 2/1/1
10.20.1.6/32 Push -- 262137
10.10.11.4 2/1/2
10.20.1.6/32 Swap 262135 262137
10.10.11.4 2/1/2
=====

```

```
No. of IPv4 Prefix Active Bindings: 11
```

```
=====
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

3ffe::a14:102/128 Push -- 262142
fe80::122 lag-1

3ffe::a14:102/128 Swap 262138 262142
fe80::122 lag-1

3ffe::a14:103/128 Pop 262142 --
-- --

3ffe::a14:104/128 Push -- 262142
fe80::114 2/1/2

3ffe::a14:104/128 Swap 262134 262142
fe80::114 2/1/2

3ffe::a14:105/128 Push -- 262142
fe80::55 2/1/1

3ffe::a14:105/128 Swap 262132 262142
fe80::55 2/1/1

3ffe::a14:106/128 Push -- 262136
fe80::114 2/1/2

3ffe::a14:106/128 Swap 262133 262136
fe80::114 2/1/2

No. of IPv6 Prefix Active Bindings: 11
=====
*A:Dut-C#

*A:Dut-C# show router ldp bindings active prefixes ipv4
=====
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 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 |        |        |
| 10.20.1.2/32 | Push  | --     | 262143 |
| 10.10.12.2   | lag-1 |        |        |
| 10.20.1.2/32 | Swap  | 262140 | 262143 |
| 10.10.12.2   | lag-1 |        |        |
| 10.20.1.3/32 | Pop   | 262143 | --     |
| --           | --    |        |        |
| 10.20.1.4/32 | Push  | --     | 262143 |
| 10.10.11.4   | 2/1/2 |        |        |
| 10.20.1.4/32 | Swap  | 262139 | 262143 |
| 10.10.11.4   | 2/1/2 |        |        |
| 10.20.1.5/32 | Push  | --     | 262143 |
| 10.10.5.5    | 2/1/1 |        |        |
| 10.20.1.5/32 | Swap  | 262137 | 262143 |
| 10.10.5.5    | 2/1/1 |        |        |
| 10.20.1.6/32 | Push  | --     | 262137 |
| 10.10.11.4   | 2/1/2 |        |        |
| 10.20.1.6/32 | Swap  | 262135 | 262137 |
| 10.10.11.4   | 2/1/2 |        |        |

-----  
No. of IPv4 Prefix Active Bindings: 11  
=====

\*A:Dut-C#

\*A:Dut-C# show router ldp bindings active prefixes ipv6

=====

|                                        |
|----------------------------------------|
| 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 IPv6 Prefix Bindings (Active)

=====

| Prefix<br>EgrNextHop | Op<br>EgrIf/LspId | IngLbl | EgrLbl |
|----------------------|-------------------|--------|--------|
| 3ffe::a14:101/128    | Push              | --     | 262142 |
| fe80::21             | 1/1/1             |        |        |
| 3ffe::a14:101/128    | Swap              | 262136 | 262142 |
| fe80::21             | 1/1/1             |        |        |
| 3ffe::a14:102/128    | Push              | --     | 262142 |

|                   |       |        |        |
|-------------------|-------|--------|--------|
| fe80::122         | lag-1 |        |        |
| 3ffe::a14:102/128 | Swap  | 262138 | 262142 |
| fe80::122         | lag-1 |        |        |
| 3ffe::a14:103/128 | Pop   | 262142 | --     |
| --                | --    |        |        |
| 3ffe::a14:104/128 | Push  | --     | 262142 |
| fe80::114         | 2/1/2 |        |        |
| 3ffe::a14:104/128 | Swap  | 262134 | 262142 |
| fe80::114         | 2/1/2 |        |        |
| 3ffe::a14:105/128 | Push  | --     | 262142 |
| fe80::55          | 2/1/1 |        |        |
| 3ffe::a14:105/128 | Swap  | 262132 | 262142 |
| fe80::55          | 2/1/1 |        |        |
| 3ffe::a14:106/128 | Push  | --     | 262136 |
| fe80::114         | 2/1/2 |        |        |
| 3ffe::a14:106/128 | Swap  | 262133 | 262136 |
| fe80::114         | 2/1/2 |        |        |

-----  
No. of IPv6 Prefix Active Bindings: 11  
=====

\*A:Dut-C#

\*A:Dut-C# show router ldp bindings active prefixes prefix 3ffe::a14:101/128 detail  
egress-if 1/1/1  
=====

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 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
=====
*A:Dut-C#

*A:Dut-C# show router ldp bindings active prefixes prefix 3ffe::a14:101/128 egress-
nh fe80::21
=====
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 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
=====
*A:Dut-C#

*A:SRU4# show router ldp bindings active p2mp
=====
LDP Bindings (IPv4 LSR ID 110.20.1.4:0)
 (IPv6 LSR ID 3ffe::6e14:104[0])
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Generic IPv4 P2MP Bindings (Active)
=====
P2MP-Id Interface
RootAddr Op IngLbl EgrLbl
EgrNH EgrIf/LspId

8193 77156
110.20.1.1 Pop 255042 --
-- --
8193 77156
110.20.1.1 Swap 255042 259773
180.4.110.110 3/1/5:1
8193 77156
110.20.1.1 Pop 258780BU --

```

```

--
8193
110.20.1.1
180.4.110.110
77156
Swap
3/1/5:1
258780BU 259773

8194
110.20.1.1
--
77157
Pop
255041 --

8194
110.20.1.1
180.4.110.110
77157
Swap
3/1/5:1
255041 259772

8194
110.20.1.1
--
77157
Pop
258779BU --

8194
110.20.1.1
180.4.110.110
77157
Swap
3/1/5:1
258779BU 259772

8195
110.20.1.1
--
77158
Pop
255040 --

8195
110.20.1.1
180.4.110.110
77158
Swap
3/1/5:1
255040 259769

8195
110.20.1.1
--
77158
Pop
258777BU --

8195
110.20.1.1
*A:SRU4#
77158
Swap
258777BU 259769

```

The following show output displays recursive FECs.

```

A:Dut-C# show router ldp bindings active p2mp
=====
LDP Bindings (IPv4 LSR ID 10.20.1.3)
(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 Recursive with In-Band-SSM IPv4 P2MP Bindings (Active)
=====
RootAddr
InnerRootAddr
Source
Group
EgrNH
Interface
Op
EgrIf/LspId
IngLbl
EgrLbl

3.4.0.2
10.20.1.6

```

```

6.0.101.10 73728
224.1.1.1 Pop 262135 --
--
3.4.0.2 (UF)
10.20.1.6
6.0.101.10 73728
224.1.1.1 Swap 262135 Stitched
--
10.20.1.3 (LF)
10.20.1.6
6.0.101.10 Unknw
224.1.1.1 Push -- 262139
2.3.0.1 lag-1
10.20.1.3 (LF)
10.20.1.6
6.0.101.10 Unknw
224.1.1.1 Push -- 262139
2.33.0.1 1/1/2:2

```

```

No. of In-Band-SSM IPv4 P2MP Active Bindings: 4
=====

```

```

*A:SRU4# show router ldp bindings active p2mp p2mp-id 1 root 110.20.1.2
=====

```

```

LDP Bindings (IPv4 LSR ID 110.20.1.4:0)
(IPv6 LSR ID 3ffe::6e14:104[0])
=====

```

```

Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====

```

```

LDP Generic IPv4 P2MP Bindings (Active)
=====

```

| P2MP-Id     | Interface   | IngLbl   | EgrLbl |
|-------------|-------------|----------|--------|
| RootAddr    | Op          |          |        |
| EgrNH       | EgrIf/LspId |          |        |
| 1           | 73728       |          |        |
| 110.20.1.2  | Pop         | 253348   | --     |
| --          | --          |          |        |
| 1           | 73728       |          |        |
| 110.20.1.2  | Swap        | 253348   | 256245 |
| 170.70.58.6 | 3/2/3:8     |          |        |
| 1           | 73728       |          |        |
| 110.20.1.2  | Pop         | 260103BU | --     |
| --          | --          |          |        |
| 1           | 73728       |          |        |
| 110.20.1.2  | Swap        | 260103BU | 256245 |
| 170.70.58.6 | 3/2/3:8     |          |        |

```

No. of Generic IPv4 P2MP Active Bindings: 4
=====

```

```

*A:SRU4#

```

```

*A:SRU4# show router ldp bindings active p2mp p2mp-id 1 root 110.20.1.2 summary
No. of Generic IPv4 P2MP Active Bindings: 4
*A:SRU4# show router ldp bindings active p2mp p2mp-id 1 root 110.20.1.2 detail
=====
LDP Bindings (IPv4 LSR ID 110.20.1.4:0)
(IPv6 LSR ID 3ffe::6e14:104[0])
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Generic IPv4 P2MP Bindings (Active)
=====

P2MP Type : 1 P2MP-Id : 1
Root-Addr : 110.20.1.2

Op : Pop
Ing Lbl : 253348
Egr Lbl : --
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None Ing. Flags : None

P2MP Type : 1 P2MP-Id : 1
Root-Addr : 110.20.1.2

Op : Swap
Ing Lbl : 253348
Egr Lbl : 256245
Egr Int/LspId : 3/2/3:8
EgrNextHop : 170.70.58.6
Egr. Flags : None Ing. Flags : None
Egr If Name : src-1.8
Metric : 1 Mtu : 1500

P2MP Type : 1 P2MP-Id : 1
Root-Addr : 110.20.1.2

Op : Pop
Ing Lbl : 260103BU
Egr Lbl : --
Egr Int/LspId : --
EgrNextHop : --
Egr. Flags : None Ing. Flags : None

P2MP Type : 1 P2MP-Id : 1
Root-Addr : 110.20.1.2

Op : Swap
Ing Lbl : 260103BU
Egr Lbl : 256245
Egr Int/LspId : 3/2/3:8
EgrNextHop : 170.70.58.6
Egr. Flags : None Ing. Flags : None
Egr If Name : src-1.8
Metric : 1 Mtu : 1500
=====
No. of Generic IPv4 P2MP Active Bindings: 4
=====

```

```
*A:SRU4#

*A:SRU4# show router ldp bindings active p2mp p2mp-
id 1 root 110.20.1.2 detail egress-
if 3/2/3:8
=====
LDP Bindings (IPv4 LSR ID 110.20.1.4:0)
 (IPv6 LSR ID 3ffe::6e14:104[0])
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Generic IPv4 P2MP Bindings (Active)
=====

P2MP Type : 1 P2MP-Id : 1
Root-Addr : 110.20.1.2

Op : Swap
Ing Lbl : 253348
Egr Lbl : 256245
Egr Int/LspId : 3/2/3:8
EgrNextHop : 170.70.58.6
Egr. Flags : None Ing. Flags : None
Egr If Name : src-1.8
Metric : 1 Mtu : 1500

P2MP Type : 1 P2MP-Id : 1
Root-Addr : 110.20.1.2

Op : Swap
Ing Lbl : 260103BU
Egr Lbl : 256245
Egr Int/LspId : 3/2/3:8
EgrNextHop : 170.70.58.6
Egr. Flags : None Ing. Flags : None
Egr If Name : src-1.8
Metric : 1 Mtu : 1500
=====
No. of Generic IPv4 P2MP Active Bindings: 2
=====
*A:SRU4#

*A:SRU4# show router ldp bindings active p2mp p2mp-id 1 root 110.20.1.2 egress-nh
170.70.58.6
=====
LDP Bindings (IPv4 LSR ID 110.20.1.4:0)
 (IPv6 LSR ID 3ffe::6e14:104[0])
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Generic IPv4 P2MP Bindings (Active)
=====

P2MP-Id Interface
RootAddr Op IngLbl EgrLbl
EgrNH EgrIf/LspId

```



```
1 73728
110.20.1.2 Swap 253348 256245
170.70.58.6 3/2/3:8

1 73728
110.20.1.2 Swap 260103BU 256245
170.70.58.6 3/2/3:8
```

-----  
No. of Generic IPv4 P2MP Active Bindings: 2  
=====

\*A:SRU4#

\*A:SRU4# show router ldp bindings active p2mp ipv4 summary

No. of Generic IPv4 P2MP Active Bindings: 8870

No. of In-Band-SSM IPv4 P2MP Active Bindings: 182

No. of In-Band-VPN-SSM IPv4 P2MP Active Bindings: 0

\*A:SRU4#

\*A:SRU4# show router ldp bindings active p2mp ipv4 detail

=====

LDP Bindings (IPv4 LSR ID 110.20.1.4:0)

(IPv6 LSR ID 3ffe::6e14:104[0])

=====

Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn

WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route

=====

LDP Generic IPv4 P2MP Bindings (Active)

=====

-----  
P2MP Type : 1 P2MP-Id : 8193  
Root-Addr : 110.20.1.1  
-----

Op : Pop  
Ing Lbl : 255042  
Egr Lbl : --  
Egr Int/LspId : --  
EgrNextHop : --  
Egr. Flags : None Ing. Flags : None  
-----

P2MP Type : 1 P2MP-Id : 8193  
Root-Addr : 110.20.1.1  
-----

Op : Swap  
Ing Lbl : 255042  
Egr Lbl : 259773  
Egr Int/LspId : 3/1/5:1  
EgrNextHop : 180.4.110.110  
Egr. Flags : None Ing. Flags : None  
Egr If Name : sicily1-1:1  
Metric : 1 Mtu : 9194  
-----

P2MP Type : 1 P2MP-Id : 8193  
Root-Addr : 110.20.1.1  
-----

Op : Pop  
Ing Lbl : 258780BU  
Egr Lbl : --  
Egr Int/LspId : --

```

EgrNextHop : --
Egr. Flags : None Ing. Flags : None

P2MP Type : 1 P2MP-Id : 8193
Root-Addr : 110.20.1.1

Op : Swap
Ing Lbl : 258780BU
Egr Lbl : 259773
Egr Int/LspId : 3/1/5:1
EgrNextHop : 180.4.110.110
Egr. Flags : None Ing. Flags : None
Egr If Name : sicily1-1:1
Metric : 1 Mtu : 9194

P2MP Type : 1 P2MP-Id : 8194
Root-Addr : 110.20.1.1

Op : Pop
Ing Lbl : 255041
*A:SRU4#

```

```

*A:SRU4# show router ldp bindings active p2mp ipv4 egress-if 3/1/5:1

```

```

=====
LDP Bindings (IPv4 LSR ID 110.20.1.4:0)
 (IPv6 LSR ID 3ffe::6e14:104[0])
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====
LDP Generic IPv4 P2MP Bindings (Active)
=====

```

| P2MP-Id       | Interface   | IngLbl   | EgrLbl |
|---------------|-------------|----------|--------|
| RootAddr      | Op          |          |        |
| EgrNH         | EgrIf/LspId |          |        |
| 8193          | 77156       |          |        |
| 110.20.1.1    | Swap        | 255042   | 259773 |
| 180.4.110.110 | 3/1/5:1     |          |        |
| 8193          | 77156       |          |        |
| 110.20.1.1    | Swap        | 258780BU | 259773 |
| 180.4.110.110 | 3/1/5:1     |          |        |
| 8194          | 77157       |          |        |
| 110.20.1.1    | Swap        | 255041   | 259772 |
| 180.4.110.110 | 3/1/5:1     |          |        |
| 8194          | 77157       |          |        |
| 110.20.1.1    | Swap        | 258779BU | 259772 |
| 180.4.110.110 | 3/1/5:1     |          |        |
| 8195          | 77158       |          |        |
| 110.20.1.1    | Swap        | 255040   | 259769 |
| 180.4.110.110 | 3/1/5:1     |          |        |
| 8195          | 77158       |          |        |
| 110.20.1.1    | Swap        | 258777BU | 259769 |

```

180.4.110.110 3/1/5:1

8196
110.20.1.1 77159
180.4.110.110 Swap 255039 259768
 3/1/5:1

8196
110.20.1.1 77159
180.4.110.110 Swap 258772BU 259768
 3/1/5:1

8197
110.20.1.1 77160
180.4.110.110 Swap 255034 259762
 3/1/5:1

8197
110.20.1.1 77160
180.4.110.110 Swap 258758BU 259762
 3/1/5:1

8198
110.20.1.1 77161
180.4.110.110 Swap 255033 259761
 3/1/5:1

8198
110.20.1.1 77161
 Swap 258755BU 259761
*A:SRU4#

```

\*A:SRU4# show router ldp bindings active p2mp ipv4 egress-nh 180.4.110.110

```

=====
LDP Bindings (IPv4 LSR ID 110.20.1.4:0)
 (IPv6 LSR ID 3ffe::6e14:104[0])
=====

```

```

Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
=====

```

LDP Generic IPv4 P2MP Bindings (Active)

```

=====
P2MP-Id Interface
RootAddr Op IngLbl EgrLbl
EgrNH EgrIf/LspId

8193
110.20.1.1 77156
180.4.110.110 Swap 255042 259773
 3/1/5:1

8193
110.20.1.1 77156
180.4.110.110 Swap 258780BU 259773
 3/1/5:1

8194
110.20.1.1 77157
180.4.110.110 Swap 255041 259772
 3/1/5:1

8194
110.20.1.1 77157
180.4.110.110 Swap 258779BU 259772
 3/1/5:1

8195
110.20.1.1 77158
 Swap 255040 259769

```

|               |         |          |        |
|---------------|---------|----------|--------|
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8195          | 77158   |          |        |
| 110.20.1.1    | Swap    | 258777BU | 259769 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8196          | 77159   |          |        |
| 110.20.1.1    | Swap    | 255039   | 259768 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8196          | 77159   |          |        |
| 110.20.1.1    | Swap    | 258772BU | 259768 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8197          | 77160   |          |        |
| 110.20.1.1    | Swap    | 255034   | 259762 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8197          | 77160   |          |        |
| 110.20.1.1    | Swap    | 258758BU | 259762 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8198          | 77161   |          |        |
| 110.20.1.1    | Swap    | 255033   | 259761 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8198          | 77161   |          |        |
| 110.20.1.1    | Swap    | 258755BU | 259761 |
| *A:SRU4#      |         |          |        |

\*A:SRU4# show router ldp bindings active p2mp ipv4 egress-if 3/1/5:1 opaque-type generic

=====

LDP Bindings (IPv4 LSR ID 110.20.1.4:0)

(IPv6 LSR ID 3ffe::6e14:104[0])

=====

Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn

WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route

=====

LDP Generic IPv4 P2MP Bindings (Active)

=====

| P2MP-Id       | Interface   |          |        |
|---------------|-------------|----------|--------|
| RootAddr      | Op          | IngLbl   | EgrLbl |
| EgrNH         | EgrIf/LspId |          |        |
| -----         |             |          |        |
| 8193          | 77156       |          |        |
| 110.20.1.1    | Swap        | 255042   | 259773 |
| 180.4.110.110 | 3/1/5:1     |          |        |
| 8193          | 77156       |          |        |
| 110.20.1.1    | Swap        | 258780BU | 259773 |
| 180.4.110.110 | 3/1/5:1     |          |        |
| 8194          | 77157       |          |        |
| 110.20.1.1    | Swap        | 255041   | 259772 |
| 180.4.110.110 | 3/1/5:1     |          |        |
| 8194          | 77157       |          |        |

|               |         |          |        |
|---------------|---------|----------|--------|
| 110.20.1.1    | Swap    | 258779BU | 259772 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8195          | 77158   |          |        |
| 110.20.1.1    | Swap    | 255040   | 259769 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8195          | 77158   |          |        |
| 110.20.1.1    | Swap    | 258777BU | 259769 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8196          | 77159   |          |        |
| 110.20.1.1    | Swap    | 255039   | 259768 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8196          | 77159   |          |        |
| 110.20.1.1    | Swap    | 258772BU | 259768 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8197          | 77160   |          |        |
| 110.20.1.1    | Swap    | 255034   | 259762 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8197          | 77160   |          |        |
| 110.20.1.1    | Swap    | 258758BU | 259762 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8198          | 77161   |          |        |
| 110.20.1.1    | Swap    | 255033   | 259761 |
| 180.4.110.110 | 3/1/5:1 |          |        |
| 8198          | 77161   |          |        |
| 110.20.1.1    | Swap    | 258755BU | 259761 |
| *A:SRU4#      |         |          |        |

The following show output displays recursive FECs.

```
A:Dut-C# show router ldp bindings active p2mp source 6.0.101.10 group 224.1.1.1
=====
LDP Bindings (IPv4 LSR ID 10.20.1.3)
 (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 Recursive with In-Band-SSM IPv4 P2MP Bindings (Active)
=====
RootAddr
InnerRootAddr
Source Interface
Group Op IngLbl EgrLbl
EgrNH EgrIf/LspId

3.4.0.2
10.20.1.6
6.0.101.10 73728
```

```

224.1.1.1 Pop 262135 --
--
3.4.0.2 (UF)
10.20.1.6
6.0.101.10 73728
224.1.1.1 Swap 262135 Stitched
--
10.20.1.3 (LF)
10.20.1.6
6.0.101.10 Unknw
224.1.1.1 Push -- 262139
2.3.0.1 lag-1
10.20.1.3 (LF)
10.20.1.6
6.0.101.10 Unknw
224.1.1.1 Push -- 262139
2.3.0.1 1/1/2:2

```

```

No. of In-Band-SSM IPv4 P2MP Active Bindings: 4
=====

```

```

A:Dut-C# show router ldp bindings active p2mp source 6.0.101.10 group 224.1.1.1
innermost-root 10.20.1.6

```

```

=====
LDP Bindings (IPv4 LSR ID 10.20.1.3)
(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 Recursive with In-Band-SSM IPv4 P2MP Bindings (Active)
=====

```

```

RootAddr
InnerRootAddr
Source
Group
EgrNH
Interface
Op
EgrIf/LspId
IngLbl
EgrLbl

3.4.0.2
10.20.1.6
6.0.101.10 73728
224.1.1.1 Pop 262135 --
--
3.4.0.2 (UF)
10.20.1.6
6.0.101.10 73728
224.1.1.1 Swap 262135 Stitched
--
10.20.1.3 (LF)
10.20.1.6
6.0.101.10 Unknw
224.1.1.1 Push -- 262139
2.3.0.1 lag-1
10.20.1.3 (LF)
10.20.1.6
6.0.101.10 Unknw
224.1.1.1 Push -- 262139

```

```
2.33.0.1 1/1/2:2

No. of In-Band-SSM IPv4 P2MP Active Bindings: 4
=====

A:Dut-C# show router ldp bindings active p2mp source 6.0.101.10 group 224.1.1.1
inner-root 10.20.1.6 root 10.20.1.3
=====
LDP Bindings (IPv4 LSR ID 10.20.1.3)
(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 Recursive with In-Band-SSM IPv4 P2MP Bindings (Active)
=====
RootAddr
InnerRootAddr
Source
Group
EgrNH
Interface
Op
EgrIf/LspId
IngLbl
EgrLbl

10.20.1.3 (LF)
10.20.1.6
6.0.101.10
224.1.1.1
2.3.0.1
10.20.1.3 (LF)
10.20.1.6
6.0.101.10
224.1.1.1
2.33.0.1

No. of In-Band-SSM IPv4 P2MP Active Bindings: 2

*A:Dut-C# show router ldp bindings active prefixes prefix 3ffe::a14:101/128
=====
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 IPv6 Prefix Bindings (Active)
=====
Prefix
EgrNextHop
Op
EgrIf/LspId
IngLbl
EgrLbl

3ffe::a14:101/128
fe80::21
1/1/1
Push
--
262142

3ffe::a14:101/128
fe80::21
1/1/1
Swap
262136
262142
```

```

No. of IPv6 Prefix Active Bindings: 2
=====

*A:Dut-C# show router ldp bindings active prefixes prefix 3ffe::a14:101/128 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 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
=====
*A:Dut-C#

*A:Dut-C> show router ldp bindings active prefixes prefix 10.20.1.4/32
=====
LDP Bindings (IPv4 LSR ID 10.20.1.3)
(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
 (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
 (C) - FEC resolved with class-based-forwarding
=====
LDP IPv4 Prefix Bindings (Active)
=====
Prefix Op IngLbl EgrLbl
EgrNextHop EgrIf/LspId

```



```
10.20.1.4/32(I) Swap 262126 474389
10.10.11.4 2/1/2
```

```

No. of IPv4 Prefix Active Bindings: 1
=====
```

```
*A:Dut-C# show router ldp bindings active prefixes prefix 10.20.1.3/32
```

```
=====
LDP Bindings (IPv4 LSR ID 10.20.1.2:0)
 (IPv6 LSR ID ::[0])
=====
```

```
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
 (S) - Static (M) - Multi-homed Secondary Support
 (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
 (C) - FEC resolved with class-based-forwarding
=====
```

```
LDP IPv4 Prefix Bindings (Active)
=====
```

| Prefix<br>EgrNextHop | Op<br>EgrIf/LspId | IngLbl | EgrLbl |
|----------------------|-------------------|--------|--------|
| 10.20.1.3/32         | Push              | --     | 262143 |
| 10.20.1.3            | LspId 5           |        |        |
| 10.20.1.3/32         | Push              | --     | 262143 |
| 10.20.1.3            | LspId 6           |        |        |
| 10.20.1.3/32         | Push              | --     | 262143 |
| 10.20.1.3            | LspId 7           |        |        |
| 10.20.1.3/32         | Push              | --     | 262143 |
| 10.20.1.3            | LspId 8           |        |        |
| 10.20.1.3/32(C)      | Swap              | 262141 | 262143 |
| 10.20.1.3            | LspId 5           |        |        |
| 10.20.1.3/32(C)      | Swap              | 262141 | 262143 |
| 10.20.1.3            | LspId 6           |        |        |
| 10.20.1.3/32(C)      | Swap              | 262141 | 262143 |
| 10.20.1.3            | LspId 7           |        |        |
| 10.20.1.3/32(C)      | Swap              | 262141 | 262143 |
| 10.20.1.3            | LspId 8           |        |        |

```

No. of IPv4 Prefix Active Bindings: 8
=====
```

```
show router ldp bindings active prefixes prefix 10.20.1.3/32 detail
```

```
=====
LDP Bindings (IPv4 LSR ID 10.20.1.2:0)
 (IPv6 LSR ID ::[0])
=====
```

Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn  
 WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route  
 (S) - Static (M) - Multi-homed Secondary Support  
 (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route  
 (C) - FEC resolved with class-based-forwarding

=====

LDP IPv4 Prefix Bindings (Active)

=====

-----

|               |                |            |          |
|---------------|----------------|------------|----------|
| Prefix        | : 10.20.1.3/32 |            |          |
| Op            | : Push         |            |          |
| Ing Lbl       | : --           | Egr Lbl    | : 262143 |
| Egr Int/LspId | : LspId 5      |            |          |
| EgrNextHop    | : 10.20.1.3    |            |          |
| Egr. Flags    | : None         | Ing. Flags | : None   |
| Lsp Name      | : B_C_5        |            |          |
| Metric        | : 1000         | Mtu        | : 1492   |

-----

|               |                |            |          |
|---------------|----------------|------------|----------|
| Prefix        | : 10.20.1.3/32 |            |          |
| Op            | : Push         |            |          |
| Ing Lbl       | : --           | Egr Lbl    | : 262143 |
| Egr Int/LspId | : LspId 6      |            |          |
| EgrNextHop    | : 10.20.1.3    |            |          |
| Egr. Flags    | : None         | Ing. Flags | : None   |
| Lsp Name      | : B_C_6        |            |          |
| Metric        | : 1000         | Mtu        | : 1492   |

-----

|               |                |            |          |
|---------------|----------------|------------|----------|
| Prefix        | : 10.20.1.3/32 |            |          |
| Op            | : Push         |            |          |
| Ing Lbl       | : --           | Egr Lbl    | : 262143 |
| Egr Int/LspId | : LspId 7      |            |          |
| EgrNextHop    | : 10.20.1.3    |            |          |
| Egr. Flags    | : None         | Ing. Flags | : None   |
| Lsp Name      | : B_C_7        |            |          |
| Metric        | : 1000         | Mtu        | : 1492   |

-----

|               |                |            |          |
|---------------|----------------|------------|----------|
| Prefix        | : 10.20.1.3/32 |            |          |
| Op            | : Push         |            |          |
| Ing Lbl       | : --           | Egr Lbl    | : 262143 |
| Egr Int/LspId | : LspId 8      |            |          |
| EgrNextHop    | : 10.20.1.3    |            |          |
| Egr. Flags    | : None         | Ing. Flags | : None   |
| Lsp Name      | : B_C_8        |            |          |
| Metric        | : 1000         | Mtu        | : 1492   |

-----

|                  |                    |            |          |
|------------------|--------------------|------------|----------|
| Prefix           | : 10.20.1.3/32 (C) |            |          |
| Op               | : Swap             |            |          |
| Ing Lbl          | : 262141           | Egr Lbl    | : 262143 |
| Egr Int/LspId    | : LspId 5          |            |          |
| EgrNextHop       | : 10.20.1.3        |            |          |
| Egr. Flags       | : None             | Ing. Flags | : None   |
| Lsp Name         | : B_C_5            |            |          |
| Metric           | : 1000             | Mtu        | : 1492   |
| CBF Default LSP: | No                 | CBF FC     | : None   |

-----

|               |                    |         |          |
|---------------|--------------------|---------|----------|
| Prefix        | : 10.20.1.3/32 (C) |         |          |
| Op            | : Swap             |         |          |
| Ing Lbl       | : 262141           | Egr Lbl | : 262143 |
| Egr Int/LspId | : LspId 6          |         |          |

```
EgrNextHop : 10.20.1.3
Egr. Flags : None Ing. Flags : None
Lsp Name : B_C_6
Metric : 1000 Mtu : 1492
CBF Default LSP: No CBF FC : None

Prefix : 10.20.1.3/32 (C)
Op : Swap
Ing Lbl : 262141 Egr Lbl : 262143
Egr Int/LspId : LspId 7
EgrNextHop : 10.20.1.3
Egr. Flags : None Ing. Flags : None
Lsp Name : B_C_7
Metric : 1000 Mtu : 1492
CBF Default LSP: Yes CBF FC : be l2 af l1 h2 ef h1 nc

Prefix : 10.20.1.3/32 (C)
Op : Swap
Ing Lbl : 262141 Egr Lbl : 262143
Egr Int/LspId : LspId 8
EgrNextHop : 10.20.1.3
Egr. Flags : None Ing. Flags : None
Lsp Name : B_C_8
Metric : 1000 Mtu : 1492
CBF Default LSP: No CBF FC : None
=====
No. of IPv4 Prefix Active Bindings: 8
=====
```

detail

|                    |                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>detail</b> [ <b>session</b> <i>ip-addr</i> [ <i>label-space</i> ]] [ <i>family</i> ]                                                                                      |
| <b>Context</b>     | show>router>ldp>bindings                                                                                                                                                     |
| <b>Description</b> | This command displays details of LDP bindings.                                                                                                                               |
| <b>Parameters</b>  | <i>family</i> — Displays either IPv4 or IPv6 LDP information.<br><b>session</b> <i>ip-addr</i> [ <i>label-space</i> ] — Specifies the IP address and label space identifier. |
| <b>Values</b>      | <ip-addr[ <i>label-spa</i> *> : ipv4-address:label-space<br>ipv6-address[ <i>label-space</i> ]<br>label-space - [0 to 65535]                                                 |

ipv4

|                |                                                                                                                                                                                                                                                                      |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>  | <b>ipv4</b> [ <b>summary</b>   <b>detail</b> ] [ <b>egress-if</b> <i>port-id</i> ]<br><b>ipv4</b> [ <b>summary</b>   <b>detail</b> ] [ <b>egress-lsp</b> <i>tunnel-id</i> ]<br><b>ipv4</b> [ <b>summary</b>   <b>detail</b> ] [ <b>egress-nh</b> <i>ip-address</i> ] |
| <b>Context</b> | show>router>ldp>bindings                                                                                                                                                                                                                                             |

---

**Description** This command display LDP active IPv4 bindings.

**Parameters** **egress-if** *port-id* — Displays LDP active bindings by matching egress-if.  
**egress-lsp** *tunnel-id* — Specifies the tunnel identifier for this egress LSP.  
**Values** 0 to 4294967295  
**egress-nh** *ip-address* — Displays LDP active bindings by matching egress-nh.  
**Values**

|              |                                                                                                   |
|--------------|---------------------------------------------------------------------------------------------------|
| ipv4-address | - a.b.c.d                                                                                         |
| ipv6-address | - x:x:x:x:x:x:x (eight 16-bit pieces)<br>x:x:x:x:x:d.d.d.d<br>x - [0 to FFFF]H<br>d - [0 to 255]D |

*detail* — Displays detailed information.  
**summary** — Displays information in a summarized format.

## ipv6

**Syntax** **ipv6** [**summary** | **detail**] [**egress-if** *port-id*]  
**ipv6** [**summary** | **detail**] [**egress-lsp** *tunnel-id*]  
**ipv6** [**summary** | **detail**] [**egress-nh** *ip-address*]

**Context** show>router>ldp>bindings

**Description** This command display LDP active IPv6 bindings.

**Parameters** **egress-if** *port-id* — Displays LDP active bindings by matching egress-if.  
**egress-lsp** *tunnel-id* — Specifies the tunnel identifier for this egress LSP.  
**Values** 0 to 4294967295  
**egress-nh** *ip-address* — Displays LDP active bindings by matching egress-nh.  
**Values**

|              |                                                                                                   |
|--------------|---------------------------------------------------------------------------------------------------|
| ipv4-address | - a.b.c.d                                                                                         |
| ipv6-address | - x:x:x:x:x:x:x (eight 16-bit pieces)<br>x:x:x:x:x:d.d.d.d<br>x - [0 to FFFF]H<br>d - [0 to 255]D |

*detail* — Displays detailed information.  
**summary** — Displays information in a summarized format.

## label-type

|                    |                                                                                                                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>label-type start-label</b> <i>start-label</i> [ <b>end-label</b> <i>end-label</i> ] <i>label-type</i> [ <b>family</b> ]                                                                                                                                                                 |
| <b>Context</b>     | show>router>ldp>bindings                                                                                                                                                                                                                                                                   |
| <b>Description</b> | This command displays LDP FEC bindings by matching labels.                                                                                                                                                                                                                                 |
| <b>Parameters</b>  | <p><i>start-label</i> — Specifies a label value to begin the display.</p> <p><b>Values</b> 16 to 1048575</p> <p><i>end-label</i> — Specifies a label value to end the display.</p> <p><b>Values</b> 17 to 1048575</p> <p><b>family</b> — Displays either IPv4 or IPv6 LDP information.</p> |

## p2mp

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>p2mp p2mp-id</b> <i>identifier</i> <b>root</b> <i>ip-address</i> [<b>session</b> <i>ip-addr</i> [<i>label-space</i>]] [<b>summary</b>   <b>detail</b>]</p> <p><b>p2mp p2mp-id</b> <i>identifier</i> <b>root</b> <i>ip-address</i> [<b>session</b> <i>ip-addr</i> [<i>label-space</i>]] [<b>summary</b>   <b>detail</b>]</p> <p><b>inner-root</b> <i>ip-address</i></p> <p><b>p2mp p2mp-id</b> <i>identifier</i> [<b>rd</b> <i>rd</i>] [<b>session</b> <i>ip-addr</i> [<i>label-space</i>]] [<b>family</b>] [<b>summary</b>   <b>detail</b>]</p> <p>[<b>innermost-root</b> <i>ip-address</i>]</p> <p><b>p2mp p2mp-id</b> <i>identifier</i> <b>root</b> <i>ip-address</i> <b>rd</b> <i>rd</i> [<b>session</b> <i>ip-addr</i> [<i>label-space</i>]] [<b>family</b>]</p> <p>[<b>summary</b>   <b>detail</b>] [<b>innermost-root</b> <i>ip-address</i>]</p> <p><b>p2mp</b> [<b>session</b> <i>ip-addr</i> [<i>label-space</i>]] [<b>family</b>] [<b>summary</b>   <b>detail</b>] [<b>opaque-type</b> <i>opaque-type</i>]</p> <p><b>p2mp source</b> <i>ip-address</i> <b>group</b> <i>mcast-address</i> <b>root</b> <i>ip-address</i> [<b>session</b> <i>ip-addr</i> [<i>label-space</i>]] [<b>family</b>] [<b>summary</b>   <b>detail</b>] <b>inner-root</b> <i>ip-address</i></p> <p><b>p2mp source</b> <i>ip-address</i> <b>group</b> <i>mcast-address</i> <b>root</b> <i>ip-address</i> [<b>rd</b> <i>rd</i>] [<b>session</b> <i>ip-addr</i> [<i>label-space</i>]] [<b>summary</b>   <b>detail</b>]</p> <p><b>p2mp source</b> <i>ip-address</i> <b>group</b> <i>mcast-address</i> [<b>session</b> <i>ip-addr</i> [<i>label-space</i>]] [<b>family</b>] [<b>summary</b>   <b>detail</b>] [<b>innermost-root</b> <i>ip-address</i>]</p> |
| <b>Context</b>     | show>router>ldp>bindings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | This command displays LDP P2MP FEC bindings.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <p><b>detail</b> — Displays detailed information.</p> <p><b>family</b> — Displays either IPv4 or IPv6 active LDP information.</p> <p><b>group</b> <i>mcast-address</i> — Displays the P2MP group multicast address bindings.</p> <p><b>inner-root</b> <i>ip-address</i> — Displays recursive FECs whose inner root address matches the specified address.</p> <p><b>innermost-root</b> <i>ip-address</i> — 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.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

**opaque-type** *opaque-type* — Specifies the type of a Multi-Point Opaque Value Element.

**Values** generic, ssm, vpn-ssm, recursive-ssm, vpn-recursive, grt-recursive

**p2mp source** *ip-address* — Displays LDP active P2MP source bindings.

**Values**

|              |                                                                                                   |
|--------------|---------------------------------------------------------------------------------------------------|
| ipv4-address | - a.b.c.d                                                                                         |
| ipv6-address | - x:x:x:x:x:x:x (eight 16-bit pieces)<br>x:x:x:x:x:d.d.d.d<br>x - [0 to FFFF]H<br>d - [0 to 255]D |

**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.

**summary** — Displays information in a summarized format.

### Output

#### Sample Output

The following show output displays recursive FECs.

```
A:Dut-C# show router ldp bindings p2mp
=====
LDP Bindings (IPv4 LSR ID 10.20.1.3)
 (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 Recursive with In-Band-SSM IPv4 P2MP Bindings
=====
RootAddr
InnerRootAddr
Source
Group
EgrNH
Peer
Interface IngLbl EgrLbl
EgrIf/LspId

3.4.0.2 (UF)
10.20.1.6
```

```

6.0.101.10
224.1.1.1 73728 262135U --
--
10.20.1.4:0
10.20.1.3 (LF)
10.20.1.6
6.0.101.10
224.1.1.1 Unknw -- 262139
2.3.0.1 lag-1
10.20.1.2:0
10.20.1.3 (LF)
10.20.1.6
6.0.101.10
224.1.1.1 Unknw -- 262139
2.33.0.1 1/1/2:2
10.20.1.2:0

No. of Recursive with In-Band-SSM IPv4 P2MP Bindings: 3
=====

A:Dut-C# show router ldp bindings p2mp source 6.0.101.10 group 224.1.1.1
=====
LDP Bindings (IPv4 LSR ID 10.20.1.3)
(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 Recursive with In-Band-SSM IPv4 P2MP Bindings
=====
RootAddr
InnerRootAddr
Source
Group
EgrNH
Peer

3.4.0.2 (UF)
10.20.1.6
6.0.101.10
224.1.1.1 73728 262135U --
--
10.20.1.4:0
10.20.1.3 (LF)
10.20.1.6
6.0.101.10
224.1.1.1 Unknw -- 262139
2.3.0.1 lag-1
10.20.1.2:0
10.20.1.3 (LF)
10.20.1.6
6.0.101.10
224.1.1.1 Unknw -- 262139
2.33.0.1 1/1/2:2
10.20.1.2:0

No. of Recursive with In-Band-SSM IPv4 P2MP Bindings: 3

```

```
=====
A:Dut-C# show router ldp bindings p2mp source 6.0.101.10 group 224.1.1.1 innermost-
root 10.20.1.6
=====
```

```
LDP Bindings (IPv4 LSR ID 10.20.1.3)
 (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 Recursive with In-Band-SSM IPv4 P2MP Bindings
=====
```

| RootAddr       | Interface   | IngLbl  | EgrLbl |
|----------------|-------------|---------|--------|
| InnerRootAddr  | EgrIf/LspId |         |        |
| Source         |             |         |        |
| Group          |             |         |        |
| EgrNH          |             |         |        |
| Peer           |             |         |        |
| -----          |             |         |        |
| 3.4.0.2 (UF)   |             |         |        |
| 10.20.1.6      |             |         |        |
| 6.0.101.10     |             |         |        |
| 224.1.1.1      | 73728       | 262135U | --     |
| --             | --          |         |        |
| 10.20.1.4:0    |             |         |        |
| 10.20.1.3 (LF) |             |         |        |
| 10.20.1.6      |             |         |        |
| 6.0.101.10     |             |         |        |
| 224.1.1.1      | Unknw       | --      | 262139 |
| 2.3.0.1        | lag-1       |         |        |
| 10.20.1.2:0    |             |         |        |
| 10.20.1.3 (LF) |             |         |        |
| 10.20.1.6      |             |         |        |
| 6.0.101.10     |             |         |        |
| 224.1.1.1      | Unknw       | --      | 262139 |
| 2.33.0.1       | 1/1/2:2     |         |        |
| 10.20.1.2:0    |             |         |        |

```

No. of Recursive with In-Band-SSM IPv4 P2MP Bindings: 3
=====
```

```
A:Dut-C# show router ldp bindings p2mp source 6.0.101.10 group 224.1.1.1 inner-
root 10.20.1.6 root 10.20.1.3
=====
```

```
LDP Bindings (IPv4 LSR ID 10.20.1.3)
 (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 Recursive with In-Band-SSM IPv4 P2MP Bindings
=====
```

| RootAddr      | Interface | IngLbl | EgrLbl |
|---------------|-----------|--------|--------|
| InnerRootAddr |           |        |        |
| Source        |           |        |        |
| Group         |           |        |        |



| EgrNH                                                   | EgrIf/LspId     |
|---------------------------------------------------------|-----------------|
| Peer                                                    |                 |
| -----                                                   |                 |
| 10.20.1.3 (LF)                                          |                 |
| 10.20.1.6                                               |                 |
| 6.0.101.10                                              |                 |
| 224.1.1.1                                               | Unknw -- 262139 |
| 2.3.0.1                                                 | lag-1           |
| 10.20.1.2:0                                             |                 |
| 10.20.1.3 (LF)                                          |                 |
| 10.20.1.6                                               |                 |
| 6.0.101.10                                              |                 |
| 224.1.1.1                                               | Unknw -- 262139 |
| 2.33.0.1                                                | 1/1/2:2         |
| 10.20.1.2:0                                             |                 |
| -----                                                   |                 |
| No. of Recursive with In-Band-SSM IPv4 P2MP Bindings: 2 |                 |
| =====                                                   |                 |

prefixes

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>prefixes</b> [ <i>family</i> ] [ <b>summary</b>   <b>detail</b> ] [ <b>egress-if</b> <i>port-id</i> ]<br><b>prefixes</b> [ <i>family</i> ] [ <b>summary</b>   <b>detail</b> ] [ <b>egress-lsp</b> <i>tunnel-id</i> ]<br><b>prefixes</b> [ <b>egress-nh</b> <i>ip-address</i> ] [ <i>family</i> ] [ <b>summary</b>   <b>detail</b> ]<br><b>prefixes prefix</b> <i>ip-prefix/ip-prefix-length</i> [ <b>summary</b>   <b>detail</b> ] [ <b>egress-if</b> <i>port-id</i> ]<br><b>prefixes prefix</b> <i>ip-prefix/ip-prefix-length</i> [ <b>summary</b>   <b>detail</b> ] [ <b>egress-lsp</b> <i>tunnel-id</i> ]<br><b>prefixes prefix</b> <i>ip-prefix/ip-prefix-length</i> [ <b>egress-nh</b> <i>ip-address</i> ] [ <b>summary</b>   <b>detail</b> ] |
| <b>Context</b>     | show>router>ldp>bindings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | This command display LDP Prefix fec bindings.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <b>prefix</b> <i>ip-prefix/ip-prefix-length</i> — Specify information for the specified IP prefix and mask length.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Values</b>      | ipv4-address - a.b.c.d<br>ipv6-address - x:x:x:x:x:x:x (eight 16-bit pieces)<br>x:x:x:x:x:x.d.d.d.d<br>x - [0 to FFFF]H<br>d - [0 to 255]D                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|                    | <i>detail</i> — Displays detailed information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|                    | <b>summary</b> — Displays information in a summarized format.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|                    | <i>family</i> — Displays either IPv4 or IPv6 active LDP information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                    | <b>egress-lsp</b> <i>tunnel-id</i> — Specifies the tunnel identifier for this egress LSP.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Values</b>      | 0 to 4294967295                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

**egress-nh** *ip-address* — Displays LDP active bindings by matching egress-nh.

**Values**

ipv4-address      - a.b.c.d  
 ipv6-address      - x:x:x:x:x:x:x (eight 16-bit pieces)  
                       x:x:x:x:x:x.d.d.d.d  
                       x - [0 to FFFF]H  
                       d - [0 to 255]D

**egress-if** *port-id* — Displays LDP active bindings by matching egress-if.

## 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 about the VC type associated with this service FEC.  
                       ethernet, vlan, mirror, frdlci, atmsdu, atmcell, atmvc, atmvp, ipipe, satop-e1, satop-t1, cesopsn, cesopsn-cas

**vc-id** *vc-id* — Displays information about the VC ID associated with this service FEC.

**saii** *global-id:prefix:ac-id* — Specifies the a SAII (Source Attachment Individual Identifier).

**Values**          <number>:<number> | <a.b.c.d>:<number>

**taii** — Specifies the TAI ID, 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 2147483647

*svc-fec-type* — Specifies the FEC type.

**Values** fec128, fec129

**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

## session

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>session</b> [ <b>family</b> ] [ <b>summary</b>   <b>detail</b> ] <i>ip-addr</i> [ <i>label-space</i> ]                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | show>router>ldp>bindings                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | This command displays LDP FEC bindings by matching peer LSR ID.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><i>detail</i> — Displays detailed information.</p> <p><b>summary</b> — Displays information in a summarized format.</p> <p><i>family</i> — Displays either IPv4 or IPv6 LDP session information.</p> <p><i>ip-addr</i> — displays configuration information about LDP sessions.</p> <p><i>label-space</i> — Specifies the label space identifier that the router is advertising on the interface.</p> <p><b>Values</b> 0 to 65535</p> |

## summary

**Syntax** **summary** [**session** *ip-addr*[*label-space*]] [**ipv4** | **ipv6**]

|                    |                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | show>router>ldp>bindings                                                                                                     |
| <b>Description</b> | This command displays a summary of LDP bindings.                                                                             |
| <b>Parameters</b>  | <b>session</b> <i>ip-addr</i> [ <i>label-space</i> ] — Specifies the IP address and label space identifier.                  |
| <b>Values</b>      | <ip-addr[ <i>label-spa</i> *> : ipv4-address:label-space<br>ipv6-address[ <i>label-space</i> ]<br>label-space - [0 to 65535] |
|                    | <b>ipv4</b> — Displays IPv4 summary bindings information.                                                                    |
|                    | <b>ipv6</b> — Displays IPv6 summary bindings information.                                                                    |

## discovery

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>discovery</b> [{ <b>peer</b> <i>ip-address</i> }]   [{ <b>interface</b> <i>ip-int-name</i> }] [ <b>state</b> <i>state</i> ] [ <b>detail</b> ] [ <b>adjacency-type</b> <i>type</i> ]<br><b>discovery</b> [ <b>state</b> <i>state</i> ] [ <b>detail</b>   <b>summary</b> ] [ <b>adjacency-type</b> <i>type</i> ] [ <b>session</b> <i>ip-addr</i> [ <i>label-space</i> ]]<br><b>discovery</b> [ <b>state</b> <i>state</i> ] [ <b>detail</b>   <b>summary</b> ] [ <b>adjacency-type</b> <i>type</i> ] [ <i>family</i> ]<br><b>discovery interface</b> <i>ip-int-name</i> [ <b>state</b> <i>state</i> ] [ <b>detail</b>   <b>summary</b> ] [ <b>session</b> <i>ip-addr</i> [ <i>label-space</i> ]]<br><b>discovery peer</b> <i>ip-address</i> [ <b>state</b> <i>state</i> ] [ <b>detail</b>   <b>summary</b> ] [ <b>session</b> <i>ip-addr</i> [ <i>label-space</i> ]] |
| <b>Context</b>     | show>router>ldp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command displays the status of the interfaces participating in LDP discovery.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <i>peer ip-address</i> — Specifies to display the IP address of the peer.<br><i>interface ip-int-name</i> — The name of an existing interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.<br><i>state state</i> — Specifies to display the current operational state of the adjacency.<br><b>Values</b> established, trying, down<br><i>detail</i> — Specifies to display detailed information.<br><i>family</i> — Displays either IPv4 or IPv6 LDP session information.<br><b>adjacency-type type</b> — Specifies to display the adjacency type.<br><b>Values</b> link, targeted                                                                                                                                                                                                     |
| <b>Output</b>      | LDP Discovery Output<br><br><a href="#">Table 44</a> describes the LDP discovery output fields.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

**Table 44 LDP Discovery Output Fields**

| Label                    | Description                                                                                   |
|--------------------------|-----------------------------------------------------------------------------------------------|
| Interface Name           | The name of the interface.                                                                    |
| 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 is targeted.                          |
| State                    | Established — The adjacency is established.<br>Trying — 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.                                       |

**Sample Output**

```

*A:Dut-A# show router ldp discovery peer
=====
LDP IPv4 Hello Adjacencies
=====
Interface Name Local Addr State
AdjType Peer Addr

N/A 10.20.1.1:0 Estab
targ 10.20.1.6:0

No. of IPv4 Hello Adjacencies: 1
=====
LDP IPv6 Hello Adjacencies
=====
Interface Name Local Addr State
AdjType Peer Addr

N/A 3ffe::a14:101[0] Estab
targ 3ffe::a14:106[0]

No. of IPv6 Hello Adjacencies: 1
=====
*A:Dut-A#

*A:Dut-A# show router ldp discovery peer 10.20.1.6
=====
LDP IPv4 Hello Adjacencies
=====

```

```

Interface Name Local Addr State
AdjType Peer Addr

N/A 10.20.1.1:0 Estab
targ 10.20.1.6:0

No. of IPv4 Hello Adjacencies: 1
=====
*A:Dut-A#

*A:Dut-A# show router ldp discovery peer 10.20.1.6 detail
=====
LDP IPv4 Hello Adjacencies
=====
Peer 10.20.1.6

Local Address : 10.20.1.1:0
Peer Address : 10.20.1.6:0
Adjacency Type : targeted State : Established
Up Time : 0d 00:02:25 Hold Time Remaining : 11
Hello Mesg Recv : 39 Hello Mesg Sent : 39
Local IP Address : 10.20.1.1
Peer IP Address : 10.20.1.6
Local Hello Timeout: 15 Remote Hello Timeout: 15
Local Cfg Seq No : 3886383873 Remote Cfg Seq No : 3487172342
Lcl IPv4 P2MP Capbl: Disabled Rem IPv4 P2MP Capbl : Enabled
Lcl IPv6 P2MP Capbl: Disabled Rem IPv6 P2MP Capbl : Enabled
Lcl IPv4 Pfx Capbl: Enabled Rem IPv4 Pfx Capbl : Enabled
Lcl IPv6 Pfx Capbl: Enabled Rem IPv6 Pfx Capbl : Enabled
=====
No. of IPv4 Hello Adjacencies: 1
=====
*A:Dut-A#

*A:Dut-A# show router ldp discovery peer detail
=====
LDP IPv4 Hello Adjacencies
=====
Peer 10.20.1.6

Local Address : 10.20.1.1:0
Peer Address : 10.20.1.6:0
Adjacency Type : targeted State : Established
Up Time : 0d 00:02:48 Hold Time Remaining : 15
Hello Mesg Recv : 46 Hello Mesg Sent : 45
Local IP Address : 10.20.1.1
Peer IP Address : 10.20.1.6
Local Hello Timeout: 15 Remote Hello Timeout: 15
Local Cfg Seq No : 3886383873 Remote Cfg Seq No : 3487172342
Lcl IPv4 P2MP Capbl: Disabled Rem IPv4 P2MP Capbl : Enabled
Lcl IPv6 P2MP Capbl: Disabled Rem IPv6 P2MP Capbl : Enabled
Lcl IPv4 Pfx Capbl: Enabled Rem IPv4 Pfx Capbl : Enabled
Lcl IPv6 Pfx Capbl: Enabled Rem IPv6 Pfx Capbl : Enabled
=====
No. of IPv4 Hello Adjacencies: 1
=====

```

```
=====
LDP IPv6 Hello Adjacencies
=====

Peer 3ffe::a14:106

Local Address : 3ffe::a14:101[0]
Peer Address : 3ffe::a14:106[0]
Adjacency Type : targeted State : Established
Up Time : 0d 00:01:03 Hold Time Remaining : 34
Hello Mesg Recv : 6 Hello Mesg Sent : 5
Local IP Address : 3ffe::a14:101
Peer IP Address : 3ffe::a14:106
Local Hello Timeout: 45 Remote Hello Timeout: 45
Local Cfg Seq No : 4281565287 Remote Cfg Seq No : 1836745726
Lcl IPv4 P2MP Capbl : Disabled Rem IPv4 P2MP Capbl : Enabled
Lcl IPv6 P2MP Capbl : Disabled Rem IPv6 P2MP Capbl : Enabled
Lcl IPv4 Pfx Capbl : Enabled Rem IPv4 Pfx Capbl : Enabled
Lcl IPv6 Pfx Capbl : Enabled Rem IPv6 Pfx Capbl : Enabled
=====
No. of IPv6 Hello Adjacencies: 1
=====
*A:Dut-A#

*A:Dut-A# show router ldp discovery adjacency-type targeted
=====
LDP IPv4 Hello Adjacencies
=====

Interface Name Local Addr State
AdjType Peer Addr

N/A 10.20.1.1:0 Estab
targ 10.20.1.6:0

No. of IPv4 Hello Adjacencies: 1
=====
LDP IPv6 Hello Adjacencies
=====

Interface Name Local Addr State
AdjType Peer Addr

N/A 3ffe::a14:101[0] Estab
targ 3ffe::a14:106[0]

No. of IPv6 Hello Adjacencies: 1
=====
*A:Dut-A#

*A:Dut-A# show router ldp discovery adjacency-type targeted ipv6 state established
=====
LDP IPv6 Hello Adjacencies
=====

Interface Name Local Addr State
AdjType Peer Addr

N/A 3ffe::a14:101[0] Estab
targ 3ffe::a14:106[0]
```

```

No. of IPv6 Hello Adjacencies: 1
=====
*A:Dut-A#

*A:Dut-A# show router ldp discovery session 10.20.1.6
=====
LDP IPv4 Hello Adjacencies
=====
Interface Name Local Addr State
AdjType Peer Addr

N/A 10.20.1.1:0 Estab
targ 10.20.1.6:0

No. of IPv4 Hello Adjacencies: 1
=====
*A:Dut-A#

*A:Dut-A# show router ldp discovery interface "ip-10.10.1.1"
=====
LDP IPv4 Hello Adjacencies
=====
Interface Name Local Addr State
AdjType Peer Addr

ip-10.10.1.1 10.20.1.1:0 Estab
link 10.20.1.2:0

No. of IPv4 Hello Adjacencies: 1
=====
LDP IPv6 Hello Adjacencies
=====
Interface Name Local Addr State
AdjType Peer Addr

ip-10.10.1.1 3ffe::a14:101[0] Estab
link 3ffe::a14:102[0]

No. of IPv6 Hello Adjacencies: 1
=====
*A:Dut-A#

*A:Dut-A# show router ldp discovery interface "ip-10.10.1.1" detail
=====
LDP IPv4 Hello Adjacencies
=====
Interface "ip-10.10.1.1"

Local Address : 10.20.1.1:0
Peer Address : 10.20.1.2:0
Adjacency Type : link State : Established
Up Time : 0d 00:26:52 Hold Time Remaining : 14
```



```

Hello Mesg Recv : 426 Hello Mesg Sent : 423
Local IP Address : 10.10.1.1
Peer IP Address : 10.10.1.2
Local Hello Timeout: 15 Remote Hello Timeout: 15
Local Cfg Seq No : 3499624168 Remote Cfg Seq No : 1622338078
Lcl IPv4 P2MP Capbl : Enabled Rem IPv4 P2MP Capbl : Enabled
Lcl IPv6 P2MP Capbl : Enabled Rem IPv6 P2MP Capbl : Enabled
Lcl IPv4 Pfx Capbl : Enabled Rem IPv4 Pfx Capbl : Enabled
Lcl IPv6 Pfx Capbl : Enabled Rem IPv6 Pfx Capbl : Enabled
=====
No. of IPv4 Hello Adjacencies: 1
=====
LDP IPv6 Hello Adjacencies
=====
Interface "ip-10.10.1.1"

Local Address : 3ffe::a14:101[0]
Peer Address : 3ffe::a14:102[0]
Adjacency Type : link State : Established
Up Time : 0d 00:26:32 Hold Time Remaining : 12
Hello Mesg Recv : 421 Hello Mesg Sent : 418
Local IP Address : fe80::11
Peer IP Address : fe80::12
Local Hello Timeout: 15 Remote Hello Timeout: 15
Local Cfg Seq No : 1658693689 Remote Cfg Seq No : 4291225243
Lcl IPv4 P2MP Capbl : Enabled Rem IPv4 P2MP Capbl : Enabled
Lcl IPv6 P2MP Capbl : Enabled Rem IPv6 P2MP Capbl : Enabled
Lcl IPv4 Pfx Capbl : Enabled Rem IPv4 Pfx Capbl : Enabled
Lcl IPv6 Pfx Capbl : Enabled Rem IPv6 Pfx Capbl : Enabled
=====
No. of IPv6 Hello Adjacencies: 1
=====
*A:Dut-A#
*A:Dut-A# show router ldp discovery interface "ip-10.10.1.1" summary
 No. of IPv4 Hello Adjacencies: 1
 No. of IPv6 Hello Adjacencies: 1
*A:Dut-A#
*A:Dut-A# show router ldp discovery interface "ip-10.10.2.1" state established
=====
LDP IPv4 Hello Adjacencies
=====
Interface Name Local Addr State
AdjType Peer Addr

ip-10.10.2.1 10.20.1.1:0 Estab
link 10.20.1.3:0

No. of IPv4 Hello Adjacencies: 1
=====
LDP IPv6 Hello Adjacencies
=====
Interface Name Local Addr State
AdjType Peer Addr

ip-10.10.2.1 3ffe::a14:101[0] Estab

```

```
link 3ffe::a14:103[0]
```

```

No. of IPv6 Hello Adjacencies: 1
```

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

```
*A:Dut-A# show router ldp discovery state established
```

```
=====
LDP IPv4 Hello Adjacencies
```

```
=====
Interface Name Local Addr State
AdjType Peer Addr

N/A 10.20.1.1:0 Estab
targ 10.20.1.6:0
ip-10.10.1.1 10.20.1.1:0 Estab
link 10.20.1.2:0
ip-10.10.2.1 10.20.1.1:0 Estab
link 10.20.1.3:0
```

```

No. of IPv4 Hello Adjacencies: 3
```

```
=====
LDP IPv6 Hello Adjacencies
```

```
=====
Interface Name Local Addr State
AdjType Peer Addr

N/A 3ffe::a14:101[0] Estab
targ 3ffe::a14:106[0]
ip-10.10.1.1 3ffe::a14:101[0] Estab
link 3ffe::a14:102[0]
ip-10.10.2.1 3ffe::a14:101[0] Estab
link 3ffe::a14:103[0]
```

```

No. of IPv6 Hello Adjacencies: 3
```

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

```
*A:Dut-A# show router ldp discovery adjacency-type link
```

```
=====
LDP IPv4 Hello Adjacencies
```

```
=====
Interface Name Local Addr State
AdjType Peer Addr

ip-10.10.1.1 10.20.1.1:0 Estab
link 10.20.1.2:0
ip-10.10.2.1 10.20.1.1:0 Estab
link 10.20.1.3:0
```

```
No. of IPv4 Hello Adjacencies: 2
=====
LDP IPv6 Hello Adjacencies
=====
Interface Name Local Addr State
AdjType Peer Addr

ip-10.10.1.1 3ffe::a14:101[0] Estab
link 3ffe::a14:102[0]
ip-10.10.2.1 3ffe::a14:101[0] Estab
link 3ffe::a14:103[0]

No. of IPv6 Hello Adjacencies: 2
=====
*A:Dut-A#

*A:Dut-A# show router ldp discovery adjacency-type link ipv6
=====
LDP IPv6 Hello Adjacencies
=====
Interface Name Local Addr State
AdjType Peer Addr

ip-10.10.1.1 3ffe::a14:101[0] Estab
link 3ffe::a14:102[0]
ip-10.10.2.1 3ffe::a14:101[0] Estab
link 3ffe::a14:103[0]

No. of IPv6 Hello Adjacencies: 2
=====
*A:Dut-A#

*A:Dut-A# show router ldp discovery session 10.20.1.2
=====
LDP IPv4 Hello Adjacencies
=====
Interface Name Local Addr State
AdjType Peer Addr

ip-10.10.1.1 10.20.1.1:0 Estab
link 10.20.1.2:0
ip-10.10.2.1 10.20.1.1:0 Estab
link 10.20.1.2:0

No. of IPv4 Hello Adjacencies: 1
=====
*A:Dut-A#
*A:Dut-A# show router ldp discovery session 10.20.1.2 summary
 No. of IPv4 Hello Adjacencies: 1
*A:Dut-A#
```

## interface

|                    |                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>interface</b> [ <i>ip-int-name</i> ] [ <b>detail</b> ] [ <i>family</i> ]<br><b>interface resource-failures</b> [ <i>family</i> ]                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | show>router>ldp                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | This command displays configuration information about LDP interfaces.                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <p><i>ip-int-name</i> — The name of an existing interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p><b>detail</b> — Displays detailed information.</p> <p><i>family</i> — Displays either IPv4 or IPv6 active LDP information.</p> <p><b>resource-failures</b> — Displays which interfaces are in overload.</p> |
| <b>Output</b>      | LDP Interface Output                                                                                                                                                                                                                                                                                                                                                                                    |

[Table 45](#) describes the LDP interface output fields.

**Table 45 LDP Interface Output Fields**

| Label        | Description                                                                                                                                                                                                                                                                                          |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface    | Specifies the interface associated with the LDP instance.                                                                                                                                                                                                                                            |
| Adm          | Up — The LDP is administratively enabled.<br>Down — The LDP is administratively disabled.                                                                                                                                                                                                            |
| Opr          | Up — The LDP is operationally enabled.<br>Down — The LDP is operationally disabled.                                                                                                                                                                                                                  |
| Hello Factor | The value by which the hello timeout should be divided to give the hello time, for example, the time interval (in s), between LDP hello messages. LDP uses hello messages to discover neighbors and to detect loss of connectivity with its neighbors.                                               |
| Hold Time    | The hello time, also known as hold time. It is the time interval (in s), that LDP waits before declaring a neighbor to be down. Hello timeout is local to the system and is sent in the hello messages to a neighbor.                                                                                |
| KA Factor    | The value by which the keepalive timeout should be divided to give the keepalive time, for example, the time interval (in s), between LDP keepalive messages. LDP keepalive messages are sent to keep the LDP session from timing out when no other LDP traffic is being sent between the neighbors. |
| KA Timeout   | The time interval (in s), that LDP waits before tearing down a session. If no LDP messages are exchanged during this time interval, the LDP session is torn down. Generally the value is configured to be 3 times the keepalive time (the time interval between successive LDP keepalive messages).  |

**Sample Output**

```
*A:Dut-A# show router ldp interface
=====
LDP Interfaces
=====
Interface Adm/Opr
Sub-Interface(s) Adm/Opr Hello Hold KA KA Transport
 Fctr Time Fctr Time Address

ip-10.10.1.1 Up/Up
 ipv4 Up/Up 3 15 3 30 System
 ipv6 Up/Up 3 15 3 30 System
ip-10.10.2.1 Up/Up
 ipv4 Up/Up 3 15 3 30 System
 ipv6 Up/Up 3 15 3 30 System

No. of Interfaces: 2
=====
*A:Dut-A#

*A:Dut-A# show router ldp interface "ip-10.10.1.1"
=====
LDP Interfaces
=====
Interface Adm/Opr
Sub-Interface(s) Adm/Opr Hello Hold KA KA Transport
 Fctr Time Fctr Time Address

ip-10.10.1.1 Up/Up
 ipv4 Up/Up 3 15 3 30 System
 ipv6 Up/Up 3 15 3 30 System

No. of Interfaces: 1
=====
*A:Dut-A#

*A:Dut-A# show router ldp interface "ip-10.10.1.1" detail
=====
LDP Interfaces
=====
Interface "ip-10.10.1.1"
=====

BASE

Admin State : Up Oper State : Up
BFD Status : Disabled

IPv4

IPv4 Admin State : Up IPv4 Oper State : Up
Last Oper Chg : 0d 00:37:59
Hold Time : 15
Oper Hold Time : 15
Keepalive Timeout : 30
Transport Addr : System
Active Adjacencies : 1
Hello Factor : 3
Keepalive Factor : 3
Last Modified : 02/27/15 23:23:19
```

```
Local LSR Type : System
Local LSR : None
IPv4 Pfx Fec Cap : Enabled
IPv4 P2MP Fec Cap : Enabled
IPv6 Pfx Fec Cap : Enabled
IPv6 P2MP Fec Cap : Enabled

IPv6

IPv6 Admin State : Up
Last Oper Chg : 0d 00:37:36
Hold Time : 15
Oper Hold Time : 15
Keepalive Timeout : 30
Transport Addr : System
Active Adjacencies : 1
Local LSR Type : System
Local LSR : None
IPv4 Pfx Fec Cap : Enabled
IPv4 P2MP Fec Cap : Enabled
IPv6 Pfx Fec Cap : Enabled
IPv6 P2MP Fec Cap : Enabled
=====
No. of Interfaces: 1
=====
*A:Dut-A#

*A:Dut-A# show router ldp interface "ip-10.10.1.1" detail ipv6
=====
LDP IPv6 Interfaces
=====
Interface "ip-10.10.1.1"
=====

BASE

Admin State : Up
BFD Status : Disabled
Oper State : Up

IPv6

IPv6 Admin State : Up
Last Oper Chg : 0d 00:37:47
Hold Time : 15
Oper Hold Time : 15
Keepalive Timeout : 30
Transport Addr : System
Active Adjacencies : 1
Local LSR Type : System
Local LSR : None
IPv4 Pfx Fec Cap : Enabled
IPv4 P2MP Fec Cap : Enabled
IPv6 Pfx Fec Cap : Enabled
IPv6 P2MP Fec Cap : Enabled
=====
No. of Interfaces: 1
=====
*A:Dut-A#

*A:Dut-A# show router ldp interface resource-failures
=====
LDP IPv4 Interface Resource Failures
=====
No Matching Entries Found
```

```
=====
LDP IPv6 Interface Resource Failures
=====
No Matching Entries Found
=====
*A:Dut-A# show router ldp interface resource-failures ipv6
=====
LDP IPv6 Interface Resource Failures
=====
No Matching Entries Found
=====
*A:Dut-A#
```

## fec-egress-stats

|                    |                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fec-egress-stats</b> [ <i>ip-prefix/mask</i> ]<br><b>fec-egress-stats active</b> [ <i>family</i> ]                                                                                                                                                                                                                                                |
| <b>Context</b>     | show>router>ldp                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command displays LDP prefix FECs egress statistics.                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <i>ip-prefix</i> — Specify information for the specified IP prefix and mask length. Host bits must be 0.<br><i>mask</i> — Specifies the 32-bit address mask used to indicate the bits of an IP address that are being used for the subnet address.<br><b>Values</b> 0 to 32<br><i>family</i> — Displays either IPv4 or IPv6 LDP session information. |

### Output

#### Sample Output

```
*A:Dut-C>config>router>ldp# show router ldp fec-egress-stats 3ffe::a14:101/128
=====
LDP IPv6 FEC Egress Statistics
=====

FEC Prefix/Mask : 3ffe::a14:101/128

Collect Stats : Enabled Accounting Plcy. : Default
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
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets : 0
FC H2
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets : 0
FC EF
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets : 0
FC H1
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets : 0
FC NC
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets : 0
=====
LDP IPv6 FEC Egress Statistics: 1
=====
*A:Dut-C>config>router>ldp#

*A:Dut-C>config>router>ldp# show router ldp fec-egress-stats active
=====
LDP IPv4 FEC Egress Statistics
=====
No Matching Entries Found
=====
LDP IPv6 FEC Egress Statistics
=====

FEC Prefix/Mask : 3ffe::a14:101/128

Collect Stats : Enabled
Admin State : Up
Accounting Plcy. : Default
FC BE
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets : 0
FC L2
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets : 0
FC AF
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets : 0
FC L1
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets : 0
FC H2
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets : 0
FC EF
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets : 0
FC H1
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets : 0
FC NC
InProf Pkts : 0
InProf Octets : 0
OutProf Pkts : 0
OutProf Octets : 0

```



```

=====
LDP IPv6 FEC Egress Statistics: 1
=====

*A:Dut-C>config>router>ldp# show router ldp fec-egress-stats active ipv6
=====
LDP IPv6 FEC Egress Statistics
=====

FEC Prefix/Mask : 3ffe::a14:101/128

Collect Stats : Enabled Accounting Plcy. : Default
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 IPv6 FEC Egress Statistics: 1
=====
*A:Dut-C>config>router>ldp#

*A:Dut-C>config>router>ldp# show router ldp statistics-summary
=====
Statistics Summary
=====
LDP FEC IPv4 Prefix Egress statistics : 0
LDP FEC IPv6 Prefix Egress statistics : 1
=====
*A:Dut-C>config>router>ldp#

```

## fec-originate

**Syntax**    **fec-originate** [*ip-prefix/mask*] [*operation-type*]

- Context** show>router>ldp
- Description** This command displays LDP static prefix FECs.
- Parameters** *ip-prefix* — Specify information for the specified IP prefix and mask length. Host bits must be 0.
- mask* — Specifies the 32-bit address mask used to indicate the bits of an IP address that are being used for the subnet address.
- Values** 0 to 32
- operation-type* — Specify the operation type to display.
- Values** pop, swap
- Output** FEC Originate Output

Table 46 describes the FEC originate parameters output fields.

**Table 46 FEC Originate Output Fields**

| Label                    | Description                                                                                                                                                                                                                 |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Prefix                   | Specifies the static prefix FEC.                                                                                                                                                                                            |
| NHType                   | Specifies the type of next-hop represented by this row entry:<br>unknown — The next-hop type has not been set.<br>IP Addr — The next-hop is an IP address.<br>pop — There is no next-hop (pop the label and route).         |
| NextHop                  | The IP address of the next-hop.                                                                                                                                                                                             |
| NHIfName                 | The name of the next-hop.                                                                                                                                                                                                   |
| IngLabel                 | Specifies the label that is advertised to the upstream peer. If this variable is set to the default value of 4294967295, the ingress label will be dynamically assigned by the label manager.                               |
| EgrLabel                 | Specifies the egress label associated with this next-hop entry. The LSR will swap the incoming label with the configured egress label. If this egress label has a value of 4294967295, the LSR will pop the incoming label. |
| OprInLbl<br>OperIngLabel | Specifies the actual or operational value of the label that was advertised to the upstream peer.                                                                                                                            |

### Sample Output

```
*A:SRU4>config>router>ldp# show router ldp fec-originate
=====
LDP Static Prefix FECs
=====
```

```

Prefix NHType NextHop IngLabel EgrLabel OperIngLabel

24.1.0.0/16 Pop n/a -- -- 0
24.1.0.1/32 Pop n/a -- -- 0
24.1.0.2/32 Pop n/a -- -- 0
24.1.0.3/32 Pop n/a -- -- 0
24.1.0.4/32 Pop n/a -- -- 0
24.1.0.5/32 Pop n/a -- -- 0
24.1.0.6/32 Pop n/a -- -- 0
24.1.0.7/32 Pop n/a -- -- 0
24.1.0.8/32 Pop n/a -- -- 0
24.1.0.9/32 Pop n/a -- -- 0
...
24.251.0.0/16 Pop n/a -- -- 0
24.252.0.0/16 Pop n/a -- -- 0
24.253.0.0/16 Pop n/a -- -- 0
24.254.0.0/16 Pop n/a -- -- 0

No. of FECs: 508
=====
*A:SRU4>config>router>ldp#

*A:Dut-C>config>router>ldp# show router ldp fec-originate 3ffe::0b0b:0101/128
=====
LDP IPv6 Static Prefix FECs
=====
Prefix NHType IngLbl EgrLbl OprInLbl
NextHop
NHIfName

3ffe::b0b:101/128 Pop -- -- 0
n/a
--

No. of IPv6 Static Prefix FECs: 1
=====
*A:Dut-C>config>router>ldp#

*A:Dut-C>config>router>ldp# show router ldp fec-originate 3ffe::0b0b:0101/128 pop
=====
LDP IPv6 Static Prefix FECs
=====
Prefix NHType IngLbl EgrLbl OprInLbl
NextHop
NHIfName

3ffe::b0b:101/128 Pop -- -- 0
n/a
--

No. of IPv6 Static Prefix FECs: 1
=====
*A:Dut-C>config>router>ldp#

*A:Dut-C>config>router>ldp# show router ldp fec-originate pop
=====
LDP IPv4 Static Prefix FECs
=====
Prefix NHType NextHop IngLbl EgrLbl OprInLbl

```

```
NHIfName

No Matching Entries Found
=====
LDP IPv6 Static Prefix FECs
=====
Prefix NHType IngLbl EgrLbl OprInLbl
NextHop
NHIfName

3ffe::b0b:101/128 Pop -- -- 0
n/a
--

No. of IPv6 Static Prefix FECs: 1
=====
*A:Dut-C>config>router>ldp# show router ldp fec-originate pop ipv6
=====
LDP IPv6 Static Prefix FECs
=====
Prefix NHType IngLbl EgrLbl OprInLbl
NextHop
NHIfName

3ffe::b0b:101/128 Pop -- -- 0
n/a
--

No. of IPv6 Static Prefix FECs: 1
=====
*A:Dut-C>config>router>ldp# show router ldp fec-originate pop ipv4
=====
LDP IPv4 Static Prefix FECs
=====
Prefix NHType NextHop IngLbl EgrLbl OprInLbl
NHIfName

No Matching Entries Found
=====
*A:Dut-C>config>router>ldp#
```

parameters

- Syntax** parameters
- Context** show>router>ldp
- Description** This command displays configuration information about LDP parameters.
- Output** LDP Parameters Output

[Table 47](#) describes the LDP parameters output fields.

**Table 47 LDP Parameters Output Fields**

| Label             | Description                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Keepalive Timeout | The time interval (in s), that LDP waits before tearing down a session. If no LDP messages are exchanged during this time interval, the LDP session is torn down. Generally the value is configured to be 3 times the keepalive time (the time interval between successive LDP keepalive messages).                                                                                                                     |
| Timeout Factor    | The value by which the keepalive timeout should be divided to give the keepalive time, for example, the time interval (in s), 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 hello time, also known as hold time. It is the time interval (in s), that LDP waits before declaring a neighbor to be down. Hello timeout is local to the system and is sent in the hello messages to a neighbor.                                                                                                                                                                                                   |
| Hello Factor      | The value by which the hello timeout should be divided to give the hello time, for example, the time interval (in s), between LDP hello messages. LDP uses hello messages to discover neighbors and to detect loss of connectivity with its neighbors.                                                                                                                                                                  |
| Auth              | Enabled — Authentication using MD5 message based digest protocol is enabled.<br>Disabled — No authentication is used.                                                                                                                                                                                                                                                                                                   |
| Admin Status      | inService — The LDP is administratively enabled.<br>outService — The LDP is administratively disabled.                                                                                                                                                                                                                                                                                                                  |
| Deaggregated FECs | False — LDP aggregates multiple prefixes into a single Forwarding Equivalence Class (FEC) and advertises a single label for the FEC. This value is only applicable to LDP interfaces and not for targeted sessions.<br>True — LDP de-aggregates prefixes into multiple FECs.                                                                                                                                            |
| Propagate Policy  | The Propagate Policy value specifies whether the LSR should generate FECs and which FECs it should generate.<br>system — LDP will distribute label bindings only for the router's system IP address.<br>interface — LDP will distribute label bindings for all LDP interfaces.<br>all — LDP will distribute label bindings for all prefixes in the routing table.<br>none — LDP will not distribute any label bindings. |
| Transport Address | interface — The interface's 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 since only one LDP session is actually set up between the two neighbors.<br>system — The system's IP address is used to set up the LDP session between neighbors.                                                                   |

**Table 47 LDP Parameters Output Fields (Continued)**

| Label             | Description (Continued)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Label-Retention   | <p>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-nexthop becomes a true next hop, the label received earlier is then used.</p> <p>conservative — Advertised label mappings are retained only if they will be 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.</p> |
| Control Mode      | <p>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.</p> <p>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.</p>                                                                                                                                                                                                                                                                                                                                                            |
| Route Preference  | The route preference assigned to LDP routes. When multiple routes are available to a destination, the route with the lowest preference will be used. This value is only applicable to LDP interfaces and not for targeted sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Loop Detection    | <p>none — Loop detection is not supported on this router. This is the only valid value since Path Vector based loop detection is not supported.</p> <p>other — Loop detection is supported but by a method other than hopCount, pathVector, or hopCountAndPathVector.</p> <p>hopCount — Loop detection is supported by hop count only.</p> <p>pathVector — Loop detection is supported by path vector only.</p> <p>hopCountAndPathVector — Loop detection is supported by both path vector and hop count.</p>                                                                                                                                                                                                        |
| Keepalive Timeout | The factor used to derive the Keepalive interval.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Keepalive Factor  | The time interval (in s), 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, for example, the time interval (in s), between LDP hello messages. LDP uses hello messages to discover neighbors and to detect loss of connectivity with its neighbors.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Auth              | <p>Enabled — Authentication using MD5 message based digest protocol is enabled.</p> <p>Disabled — No authentication is used.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Passive-Mode      | <p>true — LDP responds only when it gets a connect request from a peer and will not attempt to actively connect to its neighbors.</p> <p>false — LDP actively tries to connect to its peers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Table 47 LDP Parameters Output Fields (Continued)**

| Label             | Description (Continued)                                                          |
|-------------------|----------------------------------------------------------------------------------|
| Targeted-Sessions | true — Targeted sessions are enabled.<br>false — Targeted sessions are disabled. |

### Sample Output

```
*A:Dut-A# 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:Dut-A#
```

## targ-peer-template

**Syntax**    **targ-peer-template** [peer-template]

---

|                    |                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | show>router>ldp                                                                                                                             |
| <b>Description</b> | This command displays the configured parameters of a targeted peer-template.                                                                |
| <b>Parameters</b>  | <i>peer-template</i> — Specifies a targeted peer template on this virtual router that participates in the LDP protocol up to 32 characters. |
| <b>Output</b>      |                                                                                                                                             |

**Sample Output**

```
*A:SR1-A# show router ldp targ-peer-template
=====
LDP Peer Template
=====

Peer Template templ1

Create Time : 01/03/70 12:48:55 Last Modified : 01/04/70 04:21:15
Admin State : Up
Hello Timeout : 45 Hello Factor : 3
Hello Reduction : Disabled Hello Reduction Fa* : 3
Keepalive Timeout : 40 Keepalive Factor : 4
Tunneling : Disabled
Local LSR : None
BFD Status : Disabled

Peer Template templ2

Create Time : 01/03/70 13:14:48 Last Modified : 01/04/70 04:47:08
Admin State : Up
Hello Timeout : 45 Hello Factor : 3
Hello Reduction : Disabled Hello Reduction Fa* : 3
Keepalive Timeout : 40 Keepalive Factor : 4
Tunneling : Disabled
Local LSR : None
BFD Status : Disabled

Peer Template templ3

Create Time : 01/03/70 15:56:30 Last Modified : 01/04/70 07:28:50
Admin State : Up
Hello Timeout : 45 Hello Factor : 3
Hello Reduction : Disabled Hello Reduction Fa* : 3
Keepalive Timeout : 40 Keepalive Factor : 4
Tunneling : Disabled
Local LSR : None
BFD Status : Disabled

Peer Template templ4

Create Time : 01/03/70 17:02:12 Last Modified : 01/04/70 08:34:32
Admin State : Up
Hello Timeout : 45 Hello Factor : 3
Hello Reduction : Disabled Hello Reduction Fa* : 3
Keepalive Timeout : 40 Keepalive Factor : 4
Tunneling : Disabled
Local LSR : None
```



```
BFD Status : Disabled
=====
```

targ-peer-template-map

|                    |                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>targ-peer-template-map</b> [ <i>template-name</i> ]<br><b>targ-peer-template-map</b> [ <i>template-name</i> ] <b>peers</b> |
| <b>Context</b>     | show>router>ldp                                                                                                               |
| <b>Description</b> | This command displays targeted peer template mappings to prefix policy.                                                       |
| <b>Parameters</b>  | <i>template-name</i> — Specifies the template name, up to 32 characters.                                                      |
| <b>Output</b>      |                                                                                                                               |

Sample Output

```
*A:SR1-A# /show router ldp targ-peer-template-map
=====
LDP Peer Template Map
=====

Peer Template templ1

Peer Policy 1 : policy1

Peer Template templ2

Peer Policy 1 : policy1
Peer Policy 2 : policy2
Peer Policy 3 : policy3

Peer Template templ3

Peer Policy 1 : policy2
=====

*A:SR1-A# /show router ldp targ-peer-template-map tldp-peers
=====
LDP Peer Template Map TLDP Peers
=====

Peer Template templ1

10.0.10.1 10.0.10.2
10.0.10.3 10.0.10.4
10.0.10.5 10.0.10.6
10.0.10.7 10.0.10.8
10.0.10.9 10.0.10.10
10.0.10.11 10.0.10.12
10.0.10.13 10.0.10.14
10.0.10.15 10.0.10.16
10.0.10.17 10.0.10.18
```

```

10.0.10.19 10.0.10.20
10.0.10.21 10.0.10.22
10.0.10.23 10.0.10.24
10.0.10.25 10.0.10.26
10.0.10.27 10.0.10.28
10.0.10.29 10.0.10.30
10.0.10.31 10.0.10.32
10.0.10.33 10.0.10.34
10.0.10.35 10.0.10.36
10.0.10.37 10.0.10.38
10.0.10.39 10.0.10.40
10.0.10.41 10.0.10.42
10.0.10.43 10.0.10.44
10.0.10.45 10.0.10.46
10.0.10.47 10.0.10.48
10.0.10.49 10.0.10.50

```

```

Peer Template templ3

```

```

30.1.3.5 30.1.3.6
30.1.3.7 30.1.3.8
30.1.3.9 30.1.3.10
30.1.3.11 30.1.3.12
30.1.3.13 30.1.3.14
=====

```

## session-parameters

- Syntax** **session-parameters** [family]  
**session-parameters** *peer-ip-address*
- Context** show>router>ldp
- Description** This command displays LDP peer information.
- Parameters** *peer-ip-address* — Specify the peer IP address.  
*family* — Displays either IPv4 or IPv6 active LDP information.
- Output** [Table 48](#) describes the LDP session-parameters output.

**Table 48 LDP Session Parameters Output Fields**

| Label         | Description                                                                               |
|---------------|-------------------------------------------------------------------------------------------|
| Peer          | The IP address of the peer.                                                               |
| TTL security  | Enabled — LDP peering sessions protected.<br>Disabled — LDP peering sessions unprotected. |
| Min-TTL-Value | Displays the minimum TTL value for an incoming packet.                                    |

**Table 48 LDP Session Parameters Output Fields (Continued)**

| Label | Description (Continued)                                                                                               |
|-------|-----------------------------------------------------------------------------------------------------------------------|
| Auth  | Enabled — Authentication using MD5 message based digest protocol is enabled.<br>Disabled — No authentication is used. |

**Sample Output**

```
*A:Dut-A# show router ldp session-parameters
=====
LDP IPv4 Session Parameters
=====

Peer : 10.20.1.2

DOD : Disabled Adv Adj Addr Only : Disabled
FEC129 Cisco Inter* : Disabled
PE-ID MAC Flush In* : Disabled
Fec Limit : 0 Fec Limit Threshold: 90
Fec Limit Log Only : Disabled
Import Policies : None Export Policies : None
IPv4 Prefix Fec Cap: Enabled IPv6 Prefix Fec Cap: Disabled
P2MP Fec Cap : Enabled
Address Export : None

Peer : 10.20.1.3

DOD : Disabled Adv Adj Addr Only : Disabled
FEC129 Cisco Inter* : Disabled
PE-ID MAC Flush In* : Disabled
Fec Limit : 0 Fec Limit Threshold: 90
Fec Limit Log Only : Disabled
Import Policies : None Export Policies : None
IPv4 Prefix Fec Cap: Enabled IPv6 Prefix Fec Cap: Disabled
P2MP Fec Cap : Enabled
Address Export : None

Peer : 10.20.1.6

DOD : Disabled Adv Adj Addr Only : Disabled
FEC129 Cisco Inter* : Disabled
PE-ID MAC Flush In* : Disabled
Fec Limit : 0 Fec Limit Threshold: 90
Fec Limit Log Only : Disabled
Import Policies : None Export Policies : None
IPv4 Prefix Fec Cap: Enabled IPv6 Prefix Fec Cap: Enabled
P2MP Fec Cap : Enabled
Address Export : None
=====
No. of IPv4 Peers: 3
=====
* indicates that the corresponding row element may have been truncated.
=====
LDP IPv6 Session Parameters
```

```
=====

Peer : 3ffe::a14:102

DOD : Disabled Adv Adj Addr Only : Disabled
FEC129 Cisco Inter*: Disabled
PE-ID MAC Flush In*: Disabled
Fec Limit : 0 Fec Limit Threshold: 90
Fec Limit Log Only : Disabled
Import Policies : None Export Policies : None
IPv4 Prefix Fec Cap: Disabled IPv6 Prefix Fec Cap: Enabled
P2MP Fec Cap : Enabled
Address Export : None

Peer : 3ffe::a14:103

DOD : Disabled Adv Adj Addr Only : Disabled
FEC129 Cisco Inter*: Disabled
PE-ID MAC Flush In*: Disabled
Fec Limit : 0 Fec Limit Threshold: 90
Fec Limit Log Only : Disabled
Import Policies : None Export Policies : None
IPv4 Prefix Fec Cap: Disabled IPv6 Prefix Fec Cap: Enabled
P2MP Fec Cap : Enabled
Address Export : None
=====
No. of IPv6 Peers: 2
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-A#

*A:Dut-A# show router ldp session-parameters 3ffe::a14:103
=====
LDP IPv6 Session Parameters
=====

Peer : 3ffe::a14:103

DOD : Disabled Adv Adj Addr Only : Disabled
FEC129 Cisco Inter*: Disabled
PE-ID MAC Flush In*: Disabled
Fec Limit : 0 Fec Limit Threshold: 90
Fec Limit Log Only : Disabled
Import Policies : None Export Policies : None
IPv4 Prefix Fec Cap: Disabled IPv6 Prefix Fec Cap: Enabled
P2MP Fec Cap : Enabled
Address Export : None
=====
No. of IPv6 Peers: 1
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-A#

*A:Dut-A# show router ldp session-parameters ipv4
=====
LDP IPv4 Session Parameters
=====

Peer : 10.20.1.2
```

```

DOD : Disabled Adv Adj Addr Only : Disabled
FEC129 Cisco Inter* : Disabled
PE-ID MAC Flush In* : Disabled
Fec Limit : 0 Fec Limit Threshold: 90
Fec Limit Log Only : Disabled
Import Policies : None Export Policies : None
IPv4 Prefix Fec Cap : Enabled IPv6 Prefix Fec Cap: Disabled
P2MP Fec Cap : Enabled
Address Export : None

Peer : 10.20.1.3

DOD : Disabled Adv Adj Addr Only : Disabled
FEC129 Cisco Inter* : Disabled
PE-ID MAC Flush In* : Disabled
Fec Limit : 0 Fec Limit Threshold: 90
Fec Limit Log Only : Disabled
Import Policies : None Export Policies : None
IPv4 Prefix Fec Cap : Enabled IPv6 Prefix Fec Cap: Disabled
P2MP Fec Cap : Enabled
Address Export : None

Peer : 10.20.1.6

DOD : Disabled Adv Adj Addr Only : Disabled
FEC129 Cisco Inter* : Disabled
PE-ID MAC Flush In* : Disabled
Fec Limit : 0 Fec Limit Threshold: 90
Fec Limit Log Only : Disabled
Import Policies : None Export Policies : None
IPv4 Prefix Fec Cap : Enabled IPv6 Prefix Fec Cap: Enabled
P2MP Fec Cap : Enabled
Address Export : None
=====
No. of IPv4 Peers: 3
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-A#
```

## statistics

|                    |                                                                    |
|--------------------|--------------------------------------------------------------------|
| <b>Syntax</b>      | <b>statistics</b>                                                  |
| <b>Context</b>     | show>router>ldp                                                    |
| <b>Description</b> | This command displays LDP instance statistics related information. |
| <b>Output</b>      |                                                                    |

### Sample Output

```
*A:Dut-A# show router ldp statistics
=====
LDP Statistics for IPv4 LSR ID 10.20.1.1:0
```

```
IPv6 LSR ID 3ffe::a14:101[0]
=====
Session/Discovery

Active IPv4 Sess : 3 Active IPv6 Sess : 3
Active IPv4 LinkAdj : 2 Active IPv6 LinkAdj : 2
Active IPv4 TargAdj : 1 Active IPv6 TargAdj : 1
Active IPv4 If : 2 Inactive IPv4 If : 0
Active IPv6 If : 2 Inactive IPv6 If : 0
Active IPv4 Peers : 1 Inactive IPv4 Peers : 0
Active IPv6 Peers : 1 Inactive IPv6 Peers : 0
IPv4 Attempted Sess : 0 IPv6 Attempted Sess : 0
IPv4 OLoad If : 0 IPv4 OLoad Targ Peers : 0
IPv6 OLoad If : 0 IPv6 OLoad Targ Peers : 0

Protocol Stats

No Hello Err : 0 Param Adv Err : 0
Max PDU Err : 0 Label Range Err : 0
Bad LDP Id Err : 0 Bad PDU Len Err : 0
Bad Mesg Len Err : 0 Bad TLV Len Err : 0
Unknown TLV Err : 0 Bad Proto Ver Err : 0
Malformed TLV Err : 0 Keepalive Expired Err : 0
Shutdown Notif Sent : 0 Shutdown Notif Recv : 0

Prefixes

IPv4 Pfx FECs Sent : 10 IPv4 Pfx FECs Recv : 10
IPv6 Pfx FECs Sent : 10 IPv6 Pfx FECs Recv : 10
IPv4PfxFecOLSessSnt : 0 IPv4PfxFecOLSessRecv : 0
IPv6PfxFecOLSessSnt : 0 IPv6PfxFecOLSessRecv : 0
IPv4PfxFecInOLoad : 0 IPv6PfxFecInOLoad : 0

P2MP

IPv4 P2MP FECs Sent : 0 IPv4 P2MP FECs Recv : 0
IPv6 P2MP FECs Sent : 0 IPv6 P2MP FECs Recv : 0
IPv4P2MPFecOLSessSn : 0 IPv4P2MPFecOLSessRecv : 0
IPv6P2MPFecOLSessSn : 0 IPv6P2MPFecOLSessRecv : 0
IPv4P2MPFecInOLoad : 0 IPv6P2MPFecInOLoad : 0

Services

Svc FEC128s Sent : 0 Svc FEC128s Recv : 0
Svc FEC129s Sent : 0 Svc FEC129s Recv : 0
Svc Fec128 OLSessSn : 0 Svc Fec128 OLSessRecv : 0
Svc Fec129 OLSessSn : 0 Svc Fec129 OLSessRecv : 0
Svc Fec128 InOLoad : 0 Svc Fec129 InOLoad : 0
=====
*A:Dut-A#
```

status

|         |                 |
|---------|-----------------|
| Syntax  | status          |
| Context | show>router>ldp |

**Description** This command displays LDP status information.

**Output** LDP Status Output

[Table 49](#) describes the LDP status output fields.

**Table 49 LDP Status Output Fields**

| Label               | Description                                                                                                                                  |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| Admin State         | Up — The LDP is administratively enabled.<br>Down — The LDP is administratively disabled.                                                    |
| Oper State          | Up — The LDP is operationally enabled.<br>Down — The LDP is operationally disabled.                                                          |
| Created at          | The date and time when the LDP instance was created.                                                                                         |
| Up Time             | The time, in hundredths of seconds, that the LDP instance has been operationally up.                                                         |
| Last Change         | The date and time when the LDP instance was last modified.                                                                                   |
| Oper Down Events    | The number of times the LDP instance has gone operationally down since the instance was created.                                             |
| 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.                                                         |
| Serv FECs Sent      | The number of labels sent to the peer associated with this FEC.                                                                              |
| Serv FECs Recv      | The number of labels 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. |

**Table 49 LDP Status Output Fields (Continued)**

| Label                 | Description (Continued)                                                                                                                  |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| 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.                                       |
| Class-Forwarding      | Indicates whether or not class-based forwarding has been enabled.                                                                        |
| Malformed TLV Err     | The number of malformed TLV value fatal 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.                                         |
| Keepalive Expired Err | The number of session Keepalive timer expired errors detected for sessions associated with this LDP instance.                            |
| Shutdown Notif Recv   | The number of shutdown notifications received related to sessions associated with this LDP instance.                                     |

**Sample Output**

```

*A:Dut-A# show router ldp session
=====
LDP IPv4 Sessions
=====
Peer LDP Id Adj Type State Msg Sent Msg Recv Up Time

10.20.1.2:0 Link Established 1724 1725 0d 01:16:29
10.20.1.3:0 Link Established 1721 1723 0d 01:16:24
10.20.1.6:0 Targeted Established 1237 1238 0d 00:54:53

No. of IPv4 Sessions: 3
=====

LDP IPv6 Sessions
=====
Peer LDP Id

```



```

Adj Type State Msg Sent Msg Recv Up Time

3ffe::a14:102[0]
 Link Established 1718 1721 0d 01:16:10
3ffe::a14:103[0]
 Link Established 1718 1717 0d 01:16:10
3ffe::a14:106[0]
 Targeted Established 598 598 0d 00:53:07

No. of IPv6 Sessions: 3
=====
*A:Dut-A#
*A:Dut-A# show router ldp session 10.20.1.2

=====
LDP IPv4 Sessions
=====
Peer LDP Id Adj Type State Msg Sent Msg Recv Up Time

10.20.1.2:0 Link Established 1728 1729 0d 01:16:42

No. of IPv4 Sessions: 1
=====
*A:Dut-A# show router ldp session 3ffe::a14:106

=====
LDP IPv6 Sessions
=====
Peer LDP Id
Adj Type State Msg Sent Msg Recv Up Time

3ffe::a14:106[0]
 Targeted Established 601 602 0d 00:53:28

No. of IPv6 Sessions: 1
=====
*A:Dut-A# show router ldp session 10.20.1.2 local-addresses

=====
LDP Session Local-Addresses
=====

Session with Peer 10.20.1.2:0,
 Local 10.20.1.1:0

IPv4 Sent Addresses:

 10.10.1.1 10.10.2.1 10.20.1.1

IPv6 Sent Addresses:

 3ffe::a0a:101
 3ffe::a0a:201
 3ffe::a14:101
 fe80::11

IPv4 Recv Addresses:

 10.10.1.2 10.10.3.2 10.10.4.2 10.10.12.2

```

```
10.20.1.2

IPv6 Recv Addresses:

 3ffe::a0a:102
 3ffe::a0a:302
 3ffe::a0a:402
 3ffe::a0a:c02
 3ffe::a14:102
 fe80::12
=====
*A:Dut-A# show router ldp session 10.20.1.2 local-addresses sent
=====
LDP Session Local-Addresses
=====

Session with Peer 10.20.1.2:0,
 Local 10.20.1.1:0

IPv4 Sent Addresses:

 10.10.1.1 10.10.2.1 10.20.1.1

IPv6 Sent Addresses:

 3ffe::a0a:101
 3ffe::a0a:201
 3ffe::a14:101
 fe80::11
=====
*A:Dut-A# show router ldp session 10.20.1.2 local-addresses recv
=====
LDP Session Local-Addresses
=====

Session with Peer 10.20.1.2:0,
 Local 10.20.1.1:0

IPv4 Recv Addresses:

 10.10.1.2 10.10.3.2 10.10.4.2 10.10.12.2
 10.20.1.2

IPv6 Recv Addresses:

 3ffe::a0a:102
 3ffe::a0a:302
 3ffe::a0a:402
 3ffe::a0a:c02
 3ffe::a14:102
 fe80::12
=====
*A:Dut-A#
*A:Dut-A# show router ldp session 10.20.1.2 local-addresses recv ip-
addr 3ffe::a14:102
=====
```

```
LDP Session Local-Addresses
=====

Session with Peer 10.20.1.2:0,
 Local 10.20.1.1:0

IPv6 Recv Addresses:

 3ffe::a14:102
=====
*A:Dut-A#

*A:Dut-A# show router ldp session 10.20.1.2 link summary
No. of IPv4 Sessions: 1
*A:Dut-A#

*A:Dut-A# show router ldp session link

=====
LDP IPv4 Sessions
=====
Peer LDP Id Adj Type State Msg Sent Msg Recv Up Time

10.20.1.2:0 Link Established 1794 1796 0d 01:19:38
10.20.1.3:0 Link Established 1792 1794 0d 01:19:33

No. of IPv4 Sessions: 2
=====

=====
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
=====

*A:Dut-A# show router ldp session link summary
No. of IPv4 Sessions: 2
No. of IPv6 Sessions: 2
*A:Dut-A#

*A:Dut-A# show router ldp session state up link

=====
LDP IPv4 Sessions
=====
Peer LDP Id Adj Type State Msg Sent Msg Recv Up Time

10.20.1.2:0 Link Established 1805 1807 0d 01:20:08
10.20.1.3:0 Link Established 1803 1805 0d 01:20:03

```

```
No. of IPv4 Sessions: 2
=====

LDP IPv6 Sessions
=====
Peer LDP Id
Adj Type State Msg Sent Msg Recv Up Time

3ffe::a14:102[0]
 Link Established 1799 1803 0d 01:19:49
3ffe::a14:103[0]
 Link Established 1799 1799 0d 01:19:49

No. of IPv6 Sessions: 2
=====
*A:Dut-A#

*A:Dut-A# show router ldp session summary
 No. of IPv4 Sessions: 3
 No. of IPv6 Sessions: 3
*A:Dut-A#

*A:Dut-A# show router ldp session local-addresses ipv4

=====
LDP Session Local-Addresses
=====

Session with Peer 10.20.1.2:0,
 Local 10.20.1.1:0

IPv4 Sent Addresses:
 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:
 10.10.1.2 10.10.3.2 10.10.4.2 10.10.12.2
 10.20.1.2

Session with Peer 10.20.1.3:0,
 Local 10.20.1.1:0

IPv4 Sent Addresses:
 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:
 10.10.2.3 10.10.3.3 10.10.5.3 10.10.11.3
 10.10.12.3 10.20.1.3

Session with Peer 10.20.1.6:0,
 Local 10.20.1.1:0

IPv4 Sent Addresses:
 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:
 10.10.9.6 10.10.10.6 10.20.1.6

Session with Peer 3ffe::a14:102[0],
```

```

Local 3ffe::a14:101[0]

IPv4 Sent Addresses:
 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:
 10.10.1.2 10.10.3.2 10.10.4.2 10.10.12.2
 10.20.1.2

Session with Peer 3ffe::a14:103[0],
 Local 3ffe::a14:101[0]

IPv4 Sent Addresses:
 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:
 10.10.2.3 10.10.3.3 10.10.5.3 10.10.11.3
 10.10.12.3 10.20.1.3

Session with Peer 3ffe::a14:106[0],
 Local 3ffe::a14:101[0]

IPv4 Sent Addresses:
 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:
 10.10.9.6 10.10.10.6 10.20.1.6
=====
*A:Dut-A#

*A:Dut-A# show router ldp session 10.20.1.2 statistics
=====
LDP IPv4 Session Statistics
=====
Message Type Sent Received

Session 10.20.1.2:0

Hello 1298 1300
Keepalive 545 545
Init 1 1
Label Mapping 5 5
Label Request 0 0
Label Release 0 0
Label Withdraw 0 0
Label Abort 0 0
Notification 1 1
Address 3 3
Address Withdraw 1 1
Capability 0 0
=====
*A:Dut-A#

*A:Dut-A# show router ldp session 10.20.1.2 statistics hello
=====

```

```

LDP IPv4 Session Statistics
=====
Message Type Sent Received

Session 10.20.1.2:0

Hello 1303 1305
=====
*A:Dut-A# show router ldp session 10.20.1.2 statistics keepalive

LDP IPv4 Session Statistics
=====
Message Type Sent Received

Session 10.20.1.2:0

Keepalive 547 547
=====
*A:Dut-A#

*A:Dut-A# show router ldp status

LDP Status for IPv4 LSR ID 10.20.1.1:0
IPv6 LSR ID ::[0]
=====
Created at : 02/18/15 20:43:15
Last Change : 02/18/15 20:43:15
Admin State : Up
IPv4 Oper State : Up IPv6 Oper State : Down
IPv4 Up Time : 0d 01:33:06 IPv6 Down Time : 0d 01:33:06
IPv4 Oper Down Rea*: n/a IPv6 Oper Down Reason: systemIpDown
IPv4 Oper Down Eve*: 0 IPv6 Oper Down Events: 0
Tunn Down Damp Time: 3 sec
Label Withdraw Del*: 0 sec
Short. TTL Local : Enabled Implicit Null Label : Disabled
Import Policies : None Short. TTL Transit : Enabled
Tunl Exp Policies : None Export Policies : None
FRR : Disabled Class-Forwarding : Enabled
MP MBB Time : 3 Mcast Upstream FRR : Disabled
Aggregate Prefix : False Agg Prefix Policies : None

Capabilities

Dynamic : Enabled P2MP : Enabled
IPv4 Prefix Fec : Enabled IPv6 Prefix Fec : Enabled
Service Fec128 : Enabled Service Fec129 : Enabled
MP MBB : Enabled Overload : Enabled
=====
* indicates that the corresponding row element may have been truncated.

```

## statistics-summary

|                    |                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>statistics-summary</b> [ <b>active</b> ] [ <b>family</b> ]                                                                   |
| <b>Context</b>     | show>router>ldp                                                                                                                 |
| <b>Description</b> | This command displays LDP statistics summary information.                                                                       |
| <b>Parameters</b>  | <b>active</b> — Displays LDP statistics for only active paths.<br><b>family</b> — Displays either IPv4 or IPv6 LDP information. |
| <b>Output</b>      |                                                                                                                                 |

### Sample Output

```
*A:Dut-A# show router ldp statistics-summary
=====
Statistics Summary
=====
LDP FEC IPv4 Prefix Egress statistics : 0
LDP FEC IPv6 Prefix Egress statistics : 0
=====
*A:Dut-A#
```

## session

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>session</b> [ <i>ip-addr</i> [ <i>label-space</i> ]] <b>local-addresses</b> [ <b>sent</b>   <b>recv</b> ] <b>ip-addr</b> <i>ip-address</i><br><b>session</b> [ <i>ip-addr</i> [ <i>label-space</i> ]] [ <i>session-type</i> ] [ <b>state</b> <i>state</i> ] [ <b>summary</b>   <b>detail</b> ]<br><b>session</b> [ <i>ip-addr</i> [ <i>label-space</i> ]] <b>local-addresses</b> [ <b>sent</b>   <b>recv</b> ] [ <i>family</i> ]<br><b>session</b> [ <i>ip-addr</i> [ <i>label-space</i> ]] [ <b>sent</b>   <b>recv</b> ] <b>overload</b> [ <b>fec-type</b> <i>fec-type</i> ]<br><b>session</b> [ <b>sent</b>   <b>recv</b> ] <b>overload</b> [ <b>fec-type</b> <i>fec-type</i> ] [ <i>family</i> ]<br><b>session</b> [ <i>ip-addr</i> [ <i>label-space</i> ]] <b>statistics</b> [ <i>packet-type</i> ] [ <i>session-type</i> ]<br><b>session</b> <b>statistics</b> [ <i>packet-type</i> ] [ <i>session-type</i> ] [ <i>family</i> ]<br><b>session</b> [ <i>session-type</i> ] [ <b>state</b> <i>state</i> ] [ <b>summary</b>   <b>detail</b> ] [ <i>family</i> ] |
| <b>Context</b>     | show>router>ldp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command displays configuration information about LDP sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <i>ip-address</i> — Specify the IP address of the LDP peer.<br><i>label-space</i> — Specifies the label space identifier that the router is advertising on the interface.<br><b>Values</b> 0 to 65535<br><b>detail</b> — Displays detailed information.<br><b>statistics</b> <i>packet-type</i> — Specify the packet type.<br><b>Values</b> hello, keepalive, init, label, notification, address                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

*session-type* — Specifies to display the session type.

**Values** link, targeted, both

**Output** LDP Session Output

Table 50 describes the LDP session output fields.

**Table 50 LDP Session Output Fields**

| Label       | Description                                                                                                                                    |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| Peer LDP ID | The IP address of the LDP peer.                                                                                                                |
| Adj Type    | The adjacency type between the LDP peer and LDP session is targeted.                                                                           |
|             | Link — Specifies that this adjacency is a result of a link hello.<br>Targeted — Specifies that this adjacency is a result of a targeted hello. |
| State       | Established — The adjacency is established.                                                                                                    |
|             | Trying — The adjacency is not yet established.                                                                                                 |
| Mesg Sent   | The number of messages sent.                                                                                                                   |
| Mesg Rcvd   | The number of messages received.                                                                                                               |
| Up Time     | The amount of time the adjacency has been enabled.                                                                                             |

**Sample Output**

```
*A:Dut-A# show router ldp session
=====
LDP IPv4 Sessions
=====
Peer LDP Id Adj Type State Msg Sent Msg Recv Up Time

10.20.1.2:0 Link Established 1724 1725 0d 01:16:29
10.20.1.3:0 Link Established 1721 1723 0d 01:16:24
10.20.1.6:0 Targeted Established 1237 1238 0d 00:54:53

No. of IPv4 Sessions: 3
=====
LDP IPv6 Sessions
=====
Peer LDP Id Adj Type State Msg Sent Msg Recv Up Time

3ffe::a14:102[0] Link Established 1718 1721 0d 01:16:10
3ffe::a14:103[0] Link Established 1718 1717 0d 01:16:10
3ffe::a14:106[0] Targeted Established 598 598 0d 00:53:07

```



```

No. of IPv6 Sessions: 3
=====
*A:Dut-A#

*A:Dut-A# show router ldp session 10.20.1.2
=====
LDP IPv4 Sessions
=====
Peer LDP Id Adj Type State Msg Sent Msg Recv Up Time

10.20.1.2:0 Link Established 1728 1729 0d 01:16:42

No. of IPv4 Sessions: 1
=====
*A:Dut-A# show router ldp session 3ffe::a14:106
=====
LDP IPv6 Sessions
=====
Peer LDP Id
Adj Type State Msg Sent Msg Recv Up Time

3ffe::a14:106[0]
 Targeted Established 601 602 0d 00:53:28

No. of IPv6 Sessions: 1
=====
*A:Dut-A# show router ldp session 10.20.1.2 local-addresses
=====
LDP Session Local-Addresses
=====
Session with Peer 10.20.1.2:0,
 Local 10.20.1.1:0

IPv4 Sent Addresses:

 10.10.1.1 10.10.2.1 10.20.1.1

IPv6 Sent Addresses:

 3ffe::a0a:101
 3ffe::a0a:201
 3ffe::a14:101
 fe80::11

IPv4 Recv Addresses:

 10.10.1.2 10.10.3.2 10.10.4.2 10.10.12.2
 10.20.1.2

IPv6 Recv Addresses:

 3ffe::a0a:102
 3ffe::a0a:302
 3ffe::a0a:402
 3ffe::a0a:c02
 3ffe::a14:102
 fe80::12
=====

```

```
*A:Dut-A# show router ldp session 10.20.1.2 local-addresses sent

=====
LDP Session Local-Addresses
=====

Session with Peer 10.20.1.2:0,
 Local 10.20.1.1:0

IPv4 Sent Addresses:

 10.10.1.1 10.10.2.1 10.20.1.1

IPv6 Sent Addresses:

 3ffe::a0a:101
 3ffe::a0a:201
 3ffe::a14:101
 fe80::11
=====
*A:Dut-A# show router ldp session 10.20.1.2 local-addresses rcv
=====
LDP Session Local-Addresses
=====

Session with Peer 10.20.1.2:0,
 Local 10.20.1.1:0

IPv4 Rcv Addresses:

 10.10.1.2 10.10.3.2 10.10.4.2 10.10.12.2
 10.20.1.2

IPv6 Rcv Addresses:

 3ffe::a0a:102
 3ffe::a0a:302
 3ffe::a0a:402
 3ffe::a0a:c02
 3ffe::a14:102
 fe80::12
=====
*A:Dut-A#

*A:Dut-A# show router ldp session 10.20.1.2 local-addresses rcv ip-
addr 3ffe::a14:102
=====
LDP Session Local-Addresses
=====

Session with Peer 10.20.1.2:0,
 Local 10.20.1.1:0

IPv6 Rcv Addresses:

 3ffe::a14:102
=====
*A:Dut-A#
```

```
*A:Dut-A# show router ldp session 10.20.1.2 link summary
```

```
No. of IPv4 Sessions: 1
```

```
*A:Dut-A#
```

```
*A:Dut-A# show router ldp session link
```

```
=====
```

```
LDP IPv4 Sessions
```

```
=====
```

| Peer LDP Id | Adj Type | State       | Msg Sent | Msg Recv | Up Time     |
|-------------|----------|-------------|----------|----------|-------------|
| 10.20.1.2:0 | Link     | Established | 1794     | 1796     | 0d 01:19:38 |
| 10.20.1.3:0 | Link     | Established | 1792     | 1794     | 0d 01:19:33 |

```

```

```
No. of IPv4 Sessions: 2
```

```
=====
```

```
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
```

```
=====
```

```
*A:Dut-A# show router ldp session link summary
```

```
No. of IPv4 Sessions: 2
```

```
No. of IPv6 Sessions: 2
```

```
*A:Dut-A#
```

```
*A:Dut-A# show router ldp session state up link
```

```
=====
```

```
LDP IPv4 Sessions
```

```
=====
```

| Peer LDP Id | Adj Type | State       | Msg Sent | Msg Recv | Up Time     |
|-------------|----------|-------------|----------|----------|-------------|
| 10.20.1.2:0 | Link     | Established | 1805     | 1807     | 0d 01:20:08 |
| 10.20.1.3:0 | Link     | Established | 1803     | 1805     | 0d 01:20:03 |

```

```

```
No. of IPv4 Sessions: 2
```

```
=====
```

```
LDP IPv6 Sessions
```

```
=====
```

| Peer LDP Id      | Adj Type | State       | Msg Sent | Msg Recv | Up Time     |
|------------------|----------|-------------|----------|----------|-------------|
| 3ffe::a14:102[0] | Link     | Established | 1799     | 1803     | 0d 01:19:49 |
| 3ffe::a14:103[0] | Link     | Established | 1799     | 1799     | 0d 01:19:49 |

```

No. of IPv6 Sessions: 2
=====
```

```
*A:Dut-A#
```

```
*A:Dut-A# show router ldp session summary
```

```
No. of IPv4 Sessions: 3
```

```
No. of IPv6 Sessions: 3
```

```
*A:Dut-A#
```

```
*A:Dut-A# show router ldp session local-addresses ipv4
```

```
=====
```

```
LDP Session Local-Addresses
```

```

Session with Peer 10.20.1.2:0,
Local 10.20.1.1:0

```

```
IPv4 Sent Addresses:
```

```
10.10.1.1 10.10.2.1 10.20.1.1
```

```
IPv4 Recv Addresses:
```

```
10.10.1.2 10.10.3.2 10.10.4.2 10.10.12.2
10.20.1.2
```

```

Session with Peer 10.20.1.3:0,
Local 10.20.1.1:0

```

```
IPv4 Sent Addresses:
```

```
10.10.1.1 10.10.2.1 10.20.1.1
```

```
IPv4 Recv Addresses:
```

```
10.10.2.3 10.10.3.3 10.10.5.3 10.10.11.3
10.10.12.3 10.20.1.3
```

```

Session with Peer 10.20.1.6:0,
Local 10.20.1.1:0

```

```
IPv4 Sent Addresses:
```

```
10.10.1.1 10.10.2.1 10.20.1.1
```

```
IPv4 Recv Addresses:
```

```
10.10.9.6 10.10.10.6 10.20.1.6
```

```

Session with Peer 3ffe::a14:102[0],
Local 3ffe::a14:101[0]

```

```
IPv4 Sent Addresses:
```

```
10.10.1.1 10.10.2.1 10.20.1.1
```

```
IPv4 Recv Addresses:
```

```

 10.10.1.2 10.10.3.2 10.10.4.2 10.10.12.2
 10.20.1.2

Session with Peer 3ffe::a14:103[0],
 Local 3ffe::a14:101[0]

IPv4 Sent Addresses:

 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:

 10.10.2.3 10.10.3.3 10.10.5.3 10.10.11.3
 10.10.12.3 10.20.1.3

Session with Peer 3ffe::a14:106[0],
 Local 3ffe::a14:101[0]

IPv4 Sent Addresses:

 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:

 10.10.9.6 10.10.10.6 10.20.1.6
=====
*A:Dut-A#

*A:Dut-A# show router ldp session 10.20.1.2 statistics

=====
LDP IPv4 Session Statistics
=====
Message Type Sent Received

Session 10.20.1.2:0

Hello 1298 1300
Keepalive 545 545
Init 1 1
Label Mapping 5 5
Label Request 0 0
Label Release 0 0
Label Withdraw 0 0
Label Abort 0 0
Notification 1 1
Address 3 3
Address Withdraw 1 1
Capability 0 0
=====
*A:Dut-A#

*A:Dut-A# show router ldp session 10.20.1.2 statistics hello
=====
LDP IPv4 Session Statistics
=====
```

```

Message Type Sent Received

Session 10.20.1.2:0

Hello 1303 1305
=====
*A:Dut-A# show router ldp session 10.20.1.2 statistics keepalive
=====
LDP IPv4 Session Statistics
=====
Message Type Sent Received

Session 10.20.1.2:0

Keepalive 547 547
=====
*A:Dut-A#
```

targ-peer

**Syntax**     **targ-peer** [*ip-address*] [**detail**]  
              **targ-peer** [**detail**] *family*  
              **targ-peer resource-failures** [*family*]

**Context**    show>router>ldp

**Description**    This command displays configuration information about targeted LDP peers.

**Parameters**    *ip-address* — The IP address of the LDP peer.  
                  *family* — The type of IP family, **ipv4** or **ipv6**.  
                  **detail** — Displays detailed information.

**Output**        LDP Targeted Peer Output

[Table 51](#) describes the LDP targeted peer output.

**Table 51        LDP Targeted Peer Output Fields**

| Label | Description                                                                               |
|-------|-------------------------------------------------------------------------------------------|
| Peer  | The IP address of the peer.                                                               |
| Adm   | Up — The LDP is administratively enabled.<br>Down — The LDP is administratively disabled. |
| Opr   | Up — The LDP is operationally enabled.<br>Down — The LDP is operationally disabled.       |

**Table 51 LDP Targeted Peer Output Fields (Continued)**

| Label        | Description (Continued)                                                                                                                                                                                                                                                                                      |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hello Factor | The value by which the hello timeout should be divided to give the hello time, for example, the time interval (in s), between LDP hello messages. LDP uses hello messages to discover neighbors and to detect loss of connectivity with its neighbors.                                                       |
| Hold Time    | The hello time or hold time. The time interval (in s), that LDP waits before declaring a neighbor to be down. Hello timeout is local to the system and is sent in the hello messages to a neighbor.                                                                                                          |
| KA Factor    | The value by which the keepalive timeout is divided to calculate the keepalive time, for example, the time interval (in s), between LDP keepalive messages. LDP keepalive messages are sent to keep the LDP session from timing out when no other LDP traffic is being sent between the neighbors.           |
| KA Timeout   | The time interval (in s), that LDP waits before tearing down a session. If no LDP messages are exchanged during this time interval, the LDP session is torn down. Generally the value is configured to be 3 times the keepalive time (the time interval between successive LDP keepalive messages).          |
| Auth         | Enabled — Authentication using MD5 message based digest protocol is enabled.<br>Disabled — No authentication is used.                                                                                                                                                                                        |
| Passive Mode | The mode used to set up LDP sessions. This value is only applicable to targeted sessions and not to LDP interfaces.<br>True — LDP responds only when it gets a connect request from a peer and will not attempt to actively connect to its neighbors.<br>False — LDP actively tries to connect to its peers. |
| Auto Create  | Specifies if a targeted peer was automatically created through service manager. For an LDP interface, this value is always false.                                                                                                                                                                            |
| No. of Peers | The total number of LDP peers.                                                                                                                                                                                                                                                                               |
| Tunneling    | Enabled — Tunneling is enabled.<br>Disabled — No tunneling is used.                                                                                                                                                                                                                                          |
| LSP          | The LSP name.                                                                                                                                                                                                                                                                                                |

### Sample Output

```
*A:Dut-A# show router ldp session
=====
LDP IPv4 Sessions
=====
Peer LDP Id Adj Type State Msg Sent Msg Recv Up Time

10.20.1.2:0 Link Established 1724 1725 0d 01:16:29
10.20.1.3:0 Link Established 1721 1723 0d 01:16:24
10.20.1.6:0 Targeted Established 1237 1238 0d 00:54:53

No. of IPv4 Sessions: 3
```

```
=====
LDP IPv6 Sessions
=====
Peer LDP Id
Adj Type State Msg Sent Msg Recv Up Time

3ffe::a14:102[0]
 Link Established 1718 1721 0d 01:16:10
3ffe::a14:103[0]
 Link Established 1718 1717 0d 01:16:10
3ffe::a14:106[0]
 Targeted Established 598 598 0d 00:53:07

No. of IPv6 Sessions: 3
=====
*A:Dut-A#
*A:Dut-A# show router ldp session 10.20.1.2

=====
LDP IPv4 Sessions
=====
Peer LDP Id Adj Type State Msg Sent Msg Recv Up Time

10.20.1.2:0 Link Established 1728 1729 0d 01:16:42

No. of IPv4 Sessions: 1
=====
*A:Dut-A# show router ldp session 3ffe::a14:106

=====
LDP IPv6 Sessions
=====
Peer LDP Id
Adj Type State Msg Sent Msg Recv Up Time

3ffe::a14:106[0]
 Targeted Established 601 602 0d 00:53:28

No. of IPv6 Sessions: 1
=====
*A:Dut-A# show router ldp session 10.20.1.2 local-addresses

=====
LDP Session Local-Addresses
=====

Session with Peer 10.20.1.2:0,
 Local 10.20.1.1:0

IPv4 Sent Addresses:

 10.10.1.1 10.10.2.1 10.20.1.1

IPv6 Sent Addresses:

 3ffe::a0a:101
 3ffe::a0a:201
```



```

3ffe::a14:101
fe80::11

IPv4 Recv Addresses:

10.10.1.2 10.10.3.2 10.10.4.2 10.10.12.2
10.20.1.2

IPv6 Recv Addresses:

3ffe::a0a:102
3ffe::a0a:302
3ffe::a0a:402
3ffe::a0a:c02
3ffe::a14:102
fe80::12
=====
*A:Dut-A# show router ldp session 10.20.1.2 local-addresses sent
=====
LDP Session Local-Addresses
=====

Session with Peer 10.20.1.2:0,
 Local 10.20.1.1:0

IPv4 Sent Addresses:

10.10.1.1 10.10.2.1 10.20.1.1

IPv6 Sent Addresses:

3ffe::a0a:101
3ffe::a0a:201
3ffe::a14:101
fe80::11
=====
*A:Dut-A# show router ldp session 10.20.1.2 local-addresses recv
=====
LDP Session Local-Addresses
=====

Session with Peer 10.20.1.2:0,
 Local 10.20.1.1:0

IPv4 Recv Addresses:

10.10.1.2 10.10.3.2 10.10.4.2 10.10.12.2
10.20.1.2

IPv6 Recv Addresses:

3ffe::a0a:102
3ffe::a0a:302
3ffe::a0a:402
3ffe::a0a:c02
3ffe::a14:102
fe80::12

```

```

=====
*A:Dut-A#
*A:Dut-A# show router ldp session 10.20.1.2 local-addresses rcv ip-
addr 3ffe::a14:102

=====
LDP Session Local-Addresses
=====

Session with Peer 10.20.1.2:0,
 Local 10.20.1.1:0

IPv6 Rcv Addresses:

 3ffe::a14:102
=====
*A:Dut-A#

*A:Dut-A# show router ldp session 10.20.1.2 link summary
No. of IPv4 Sessions: 1
*A:Dut-A#

*A:Dut-A# show router ldp session link

=====
LDP IPv4 Sessions
=====

Peer LDP Id Adj Type State Msg Sent Msg Rcv Up Time

10.20.1.2:0 Link Established 1794 1796 0d 01:19:38
10.20.1.3:0 Link Established 1792 1794 0d 01:19:33

No. of IPv4 Sessions: 2
=====

=====
LDP IPv6 Sessions
=====

Peer LDP Id Adj Type State Msg Sent Msg Rcv 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
=====
*A:Dut-A# show router ldp session link summary
No. of IPv4 Sessions: 2
No. of IPv6 Sessions: 2
*A:Dut-A#

*A:Dut-A# show router ldp session state up link

=====
LDP IPv4 Sessions

```

```
=====
Peer LDP Id Adj Type State Msg Sent Msg Recv Up Time

10.20.1.2:0 Link Established 1805 1807 0d 01:20:08
10.20.1.3:0 Link Established 1803 1805 0d 01:20:03

No. of IPv4 Sessions: 2
=====

LDP IPv6 Sessions
=====
Peer LDP Id Adj Type State Msg Sent Msg Recv Up Time

3ffe::a14:102[0] Link Established 1799 1803 0d 01:19:49
3ffe::a14:103[0] Link Established 1799 1799 0d 01:19:49

No. of IPv6 Sessions: 2
=====

*A:Dut-A#

*A:Dut-A# show router ldp session summary
No. of IPv4 Sessions: 3
No. of IPv6 Sessions: 3
*A:Dut-A#

*A:Dut-A# show router ldp session local-addresses ipv4

=====
LDP Session Local-Addresses
=====

Session with Peer 10.20.1.2:0,
Local 10.20.1.1:0

IPv4 Sent Addresses:

 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:

 10.10.1.2 10.10.3.2 10.10.4.2 10.10.12.2
 10.20.1.2

Session with Peer 10.20.1.3:0,
Local 10.20.1.1:0

IPv4 Sent Addresses:

 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:

 10.10.2.3 10.10.3.3 10.10.5.3 10.10.11.3
 10.10.12.3 10.20.1.3

```

```

Session with Peer 10.20.1.6:0,
 Local 10.20.1.1:0

IPv4 Sent Addresses:

 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:

 10.10.9.6 10.10.10.6 10.20.1.6

Session with Peer 3ffe::a14:102[0],
 Local 3ffe::a14:101[0]

IPv4 Sent Addresses:

 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:

 10.10.1.2 10.10.3.2 10.10.4.2 10.10.12.2
 10.20.1.2

Session with Peer 3ffe::a14:103[0],
 Local 3ffe::a14:101[0]

IPv4 Sent Addresses:

 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:

 10.10.2.3 10.10.3.3 10.10.5.3 10.10.11.3
 10.10.12.3 10.20.1.3

Session with Peer 3ffe::a14:106[0],
 Local 3ffe::a14:101[0]

IPv4 Sent Addresses:

 10.10.1.1 10.10.2.1 10.20.1.1

IPv4 Recv Addresses:

 10.10.9.6 10.10.10.6 10.20.1.6
=====
*A:Dut-A#

*A:Dut-A# show router ldp session 10.20.1.2 statistics

=====
LDP IPv4 Session Statistics
=====
Message Type Sent Received

Session 10.20.1.2:0

```

```

Hello 1298 1300
Keepalive 545 545
Init 1 1
Label Mapping 5 5
Label Request 0 0
Label Release 0 0
Label Withdraw 0 0
Label Abort 0 0
Notification 1 1
Address 3 3
Address Withdraw 1 1
Capability 0 0
=====
*A:Dut-A#
*A:Dut-A# show router ldp session 10.20.1.2 statistics hello

=====
LDP IPv4 Session Statistics
=====
Message Type Sent Received

Session 10.20.1.2:0

Hello 1303 1305
=====
*A:Dut-A# show router ldp session 10.20.1.2 statistics keepalive

=====
LDP IPv4 Session Statistics
=====
Message Type Sent Received

Session 10.20.1.2:0

Keepalive 547 547
=====
*A:Dut-A#
```

## tcp-session-parameters

|                    |                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>tcp-session-parameters</b><br><b>tcp-session-parameters</b> [family]<br><b>tcp-session-parameters</b> [keychain keychain]<br><b>tcp-session-parameters</b> [transport-peer-ip-address] |
| <b>Context</b>     | show>router>ldp                                                                                                                                                                           |
| <b>Description</b> | This command displays information about the TCP transport session of an LDP peer.                                                                                                         |

**Parameters**    *family* — Specifies the family type.

**Values**    ipv4, ipv6

**keychain** *keychain* — Specifies the authentication keychain name up to 32 characters in length.

*transport-peer-ip-address* — Specifies the source of the transport address.

**Values**    <transport-peer-ip\*> : ipv4-address - a.b.c.d  
                                   ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit pieces)  
                                   x:x:x:x:x:x.d.d.d.d  
                                   x - [0 to FFFF]H  
                                   d - [0 to 255]D

**Output****Sample Output**

```
*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
=====

*A:Dut-A# show router ldp tcp-session-parameters ipv4
=====
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
=====

*A:Dut-A# show router ldp tcp-session-parameters ipv6
=====
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
=====

*A:Dut-A# show router ldp tcp-session-parameters keychain "LdpAuth"
=====
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
=====
*A:Dut-A# show router ldp tcp-session-parameters
- tcp-session-parameters [family]
- tcp-session-parameters [keychain <keychain>]
- tcp-session-parameters [<transport-peer-ip-address>]

<transport-peer-ip*> : ipv4-address - a.b.c.d
 ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit
 pieces)
 x:x:x:x:x:x:d.d.d.d
 x - [0..FFFF]H
 d - [0..255]D

<family> : ipv4 | ipv6
<keychain> : auth-keychain name [32 char max]

*A:Dut-A# show router ldp tcp-session-parameters 3ffe::a14:102
=====
LDP IPv6 TCP Session Parameters
=====

Peer Transport: 3ffe::a14:102

Authentication Key : Disabled Path MTU Discovery : Disabled
Auth key chain : LdpAuth Min-TTL : 0
=====
No. of IPv6 Peers: 1
=====
*A:Dut-A#

```

### 5.32.2.2 Clear Commands

#### fec-egress-statistics

|                    |                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>fec-egress-statistics</b> [ <i>ip-prefix/mask</i> ]                                                                                                                                                                                                                                          |
| <b>Context</b>     | clear>router>ldp                                                                                                                                                                                                                                                                                |
| <b>Description</b> | This command clears LDP FEC egress statistics.                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <p><i>ip-prefix</i> — Specify information for the specified IP prefix and mask length. Host bits must be 0.</p> <p><i>mask</i> — Specifies the 32-bit address mask used to indicate the bits of an IP address that are being used for the subnet address.</p> <p><b>Values</b>      0 to 32</p> |



## instance

|                    |                                       |
|--------------------|---------------------------------------|
| <b>Syntax</b>      | <b>instance</b> [ipv4   ipv6]         |
| <b>Context</b>     | clear>router>ldp                      |
| <b>Description</b> | This command resets the LDP instance. |

## interface

|                    |                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>interface</b> [ip-int-name] [statistics] [ipv4   ipv6]                                                                                                                                                                                                    |
| <b>Context</b>     | clear>router>ldp                                                                                                                                                                                                                                             |
| <b>Description</b> | This command restarts or clears statistics for LDP interfaces.                                                                                                                                                                                               |
| <b>Parameters</b>  | <p><i>ip-int-name</i> — The name of an existing interface. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.</p> <p><i>family</i> — Specifies to clear IPv4 or IPv6 information.</p> |

### Output

#### Sample Output

```
*A:Dut-A# clear router ldp interface "ip-10.10.1.1" ipv4
*A:Dut-A#

*A:Dut-A# clear router ldp interface "ip-10.10.1.1" ipv6
*A:Dut-A#
```

## resource-failures

|                    |                                                                   |
|--------------------|-------------------------------------------------------------------|
| <b>Syntax</b>      | <b>resource-failures</b>                                          |
| <b>Context</b>     | clear>router>ldp                                                  |
| <b>Description</b> | This command clears resource overload status in the LDP instance. |

## peer

|                    |                                                                    |
|--------------------|--------------------------------------------------------------------|
| <b>Syntax</b>      | <b>peer</b> [ip-address] [statistics]                              |
| <b>Context</b>     | clear>router>ldp                                                   |
| <b>Description</b> | This command restarts or clears statistics for LDP targeted peers. |
| <b>Parameters</b>  | <i>ip-address</i> — The IP address of a targeted peer.             |

**statistics** — Clears only the statistics for a targeted peer.

## session

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>session</b> <i>ip-addr</i> [ <i>label-space</i> ] [ <b>statistics</b> ]<br><b>session</b> <i>ip-addr</i> [ <i>label-space</i> ] <b>overload</b> [ <b>fec-type</b> <b>p2mp</b>   <b>prefixes</b> <b>sub-type</b> <i>sub-type</i> ]<br><b>session</b> <i>ip-addr</i> [ <i>label-space</i> ] <b>overload</b> [ <b>fec-type</b> <b>svc-fec128</b>   <b>svc-fec129</b> ]                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | clear>router>ldp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command restarts or clears statistics for LDP sessions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <i>ip-address</i> — Clears the IP address of the session<br><i>label-space</i> — Specifies the label space identifier that the router is advertising on the interface.<br><b>Values</b> 0 to 65535<br><b>statistics</b> — Clears only the statistics for a session.<br><b>overload</b> — Clears overload information.<br><b>fec-type</b> — Clears the specified FEC type.<br><b>Values</b> p2mp, svc-fec128, svc-fec129<br><b>session</b> <i>ip-addr</i> [ <i>label-space</i> ] — Specifies the IP address and label space identifier.<br><b>Values</b> <ip-addr[ <i>label-spa</i> *> : ipv4-address:label-space<br>ipv6-address [ <i>label-space</i> ]<br>label-space — [0 to 65535] |

## statistics

|                    |                                              |
|--------------------|----------------------------------------------|
| <b>Syntax</b>      | <b>statistics</b>                            |
| <b>Context</b>     | clear>router>ldp                             |
| <b>Description</b> | This command clears LDP instance statistics. |

### 5.32.2.3 Debug Commands

The following output shows debug LDP configurations discussed in this section.

```
A:ALA-12# debug router ldp peer 10.10.10.104
A:ALA-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
A:ALA-12>debug>router>ldp#

```

## ldp

|                    |                                              |
|--------------------|----------------------------------------------|
| <b>Syntax</b>      | <b>[no] ldp</b>                              |
| <b>Context</b>     | debug>router                                 |
| <b>Description</b> | Use this command to configure LDP debugging. |

## interface

|                    |                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] interface</b> <i>interface-name</i> <i>family</i>                                                |
| <b>Context</b>     | debug>router>ldp                                                                                         |
| <b>Description</b> | Use this command for debugging an LDP interface.                                                         |
| <b>Parameters</b>  | <i>interface-name</i> — The name of an existing interface.<br><i>family</i> — Specifies the family type. |
| <b>Values</b>      | ipv4, ipv6                                                                                               |

## peer

|                    |                                                     |
|--------------------|-----------------------------------------------------|
| <b>Syntax</b>      | <b>[no] peer</b> <i>ip-address</i>                  |
| <b>Context</b>     | debug>router>ldp                                    |
| <b>Description</b> | Use this command for debugging an LDP peer.         |
| <b>Parameters</b>  | <b>ip-address</b> — The IP address of the LDP peer. |

## event

|                    |                                                            |
|--------------------|------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] event</b>                                          |
| <b>Context</b>     | debug>router>ldp>if<br>debug>router>ldp>peer               |
| <b>Description</b> | This command configures debugging for specific LDP events. |

## bindings

|                    |                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] bindings</b>                                                                                                                                                                            |
| <b>Context</b>     | debug>router>ldp>peer>event                                                                                                                                                                     |
| <b>Description</b> | This command displays debugging information about addresses and label bindings learned from LDP peers for LDP bindings.<br><br>The <b>no</b> form of the command disables the debugging output. |

## messages

|                    |                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] messages</b>                                                                                                                                                                                                                             |
| <b>Context</b>     | debug>router>ldp>if>event<br>debug>router>ldp>peer>event                                                                                                                                                                                         |
| <b>Description</b> | This command displays specific information (for example, message type, source, and destination) regarding LDP messages sent to and received from LDP peers.<br><br>The <b>no</b> form of the command disables debugging output for LDP messages. |

## packet

|                    |                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>packet [detail]<br/>no packet</b>                                                                                             |
| <b>Context</b>     | debug>router>ldp>if<br>debug>router>ldp>peer                                                                                     |
| <b>Description</b> | This command enables debugging for specific LDP packets.<br><br>The <b>no</b> form of the command disables the debugging output. |
| <b>Parameters</b>  | <b>detail</b> — Displays detailed information.                                                                                   |

## hello

|                    |                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>hello [detail]</b><br><b>no hello</b>                                                                                      |
| <b>Context</b>     | debug>router>ldp>if>packet<br>debug>router>ldp>peer>packet                                                                    |
| <b>Description</b> | This command enables debugging for LDP hello packets.<br><br>The <b>no</b> form of the command disables the debugging output. |
| <b>Parameters</b>  | <b>detail</b> — Displays detailed information.                                                                                |

## init

|                    |                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>init [detail]</b><br><b>no init</b>                                                                                       |
| <b>Context</b>     | debug>router>ldp>peer>packet                                                                                                 |
| <b>Description</b> | This command enables debugging for LDP Init packets.<br><br>The <b>no</b> form of the command disables the debugging output. |
| <b>Parameters</b>  | <b>detail</b> — Displays detailed information.                                                                               |

## keepalive

|                    |                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] keepalive</b>                                                                                                             |
| <b>Context</b>     | debug>router>ldp>peer>packet                                                                                                      |
| <b>Description</b> | This command enables debugging for LDP Keepalive packets.<br><br>The <b>no</b> form of the command disables the debugging output. |

## label

|                    |                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>label [detail]</b><br><b>no label</b>                                                                                      |
| <b>Context</b>     | debug>router>ldp>peer>packet                                                                                                  |
| <b>Description</b> | This command enables debugging for LDP Label packets.<br><br>The <b>no</b> form of the command disables the debugging output. |

**Parameters**     **detail** — Displays detailed information.

### 5.32.2.4 Tools Commands

#### fec

**Syntax**     **fec** **vc-type** *vc-type* **agi** *agi*  
**fec** **p2mp-id** *identifier* **root** *ip-address*  
**fec** **prefix** *ip-address[/mask]*  
**fec** **root** *ip-address* **source** *ip-address* **group** *mcast-address* [**rd** *rd*]  
**fec** **root** *ip-address* **source** *ip-address* **group** *mcast-address* **inner-root** *ip-address*  
**fec** **vc-type** *vc-type* **vc-id** *vc-id*  
**fec** **vc-type** *vc-type* **agi** *agi* **saii-type2** *global-id:prefix:ac-id* **taii-type2** *global-id:prefix:ac-id*

**Context**     tools>dump>router>ldp

**Description**     This command dumps information for an LDP FEC.

**Parameters**     **p2mp-id** *identifier* — Dumps LDP active P2MP identifier bindings information.  
                          **Values**     0 to 4294967295

**inner-root** *ip-address* — Dumps inner root IP address information.

**root** *ip-address* — Dumps root IP address information.

**prefix** *ip-address[/mask]* — Dumps LDP active prefix and mask information.  
                          **Values**     *ip-address[/mask]>* : ipv4-prefix a.b.c.d  
                                         ipv4-prefix-le [0 to 32]  
                                         ipv6-prefix x:x:x:x:x:x:x (eight 16-bit pieces)  
                                         x:x:x:x:x:x:d.d.d.d  
                                         x - [0..FFFF]H  
                                         d - [0..255]D  
                                         ipv6-prefix-le [0 to 128]

**source** *ip-address* — Dumps source IP address information.

**group** *mcast-address* — Dumps the group multicast address bindings.

**rd** *rd* — Dumps information for the route distinguisher.  
                          **Values**     *ip-addr:comm-val* | *2byte-asnumber:ext-comm-val* | *4byte-asnumber:comm-val*

**vc-type** *vc-type* — Dumps information for the specified VC type.

**Values** ethernet, vlan, framerelay, atm-all5, atm-cell, hdlc, ppp, cem, atm-vcc, atm-vpc, ipipe, atm-vcc-1-1, atm-vpc-1-1, atm-aal5-pdu, fr, cep, e1-satop, t1-satop, e3-satop, t3-satop, cesopsn, tdmop, cesopsn-cas, tdmop-cas, fr-dlci, mirror

**vc-id** *vc-id* — Dumps information for the specified VC-ID.

**Values** 1 to 4294967295

**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

**saii-type2** *global-id:prefix:ac-id* — Dumps Source Attachment Individual Identifier (SAII) information.

**Values** <number>:<number> | <a.b.c.d>:<number>

**taii-type2** *global-id:prefix:ac-id* — Dumps Target Attachment Individual Identifier (TAII)

**svc-fec-type** — Specifies the FEC type.

**Values** fec128, fec129

## Output

### Sample Output

```
A:Dut-A# tools dump router ldp fec root 10.20.1.3 source 6.0.101.10 group
224.1.1.1 inner-root 10.20.1.6
P2MP: root: 10.20.1.3, T: 7, L: 21 (InnerRoot: 10.20.1.6 T: 3, L:8, Src:
6.0.101.10, Grp: 224.1.1.1)
 Create Time : 01/27/16 16:39:04.097 (elapsed: 0d 03:20:24)
 Last Mod. Time: 01/27/16 16:39:04.097 (elapsed: 0d 03:20:24)
 FEC Flags : Pop UprStitched
 TunlIfId: 73728 (OperState : up)
 LSP ID : 0
 LSP ID Acct. : 0
 isIngressMttm : No HasLeaf : Yes
 isIngrIntermdte: No
 CanProgIngress: No
 InPhopFrr : No
 isStitchedUpr : Yes
 RslvdPhop(p) : 10.20.1.2:0 (seqNum 2)
 RslvdPhop(b) : 0.0.0.0:0 (seqNum 0)
```

```

pri Upstream : 10.20.1.2:0, AdvLabel 262139
mbb Upstream : None
bkb Upstream : None
AdvInLabel(p) : 262139
AdvInLabel(b) : 0
PrgInLabel(p) : 1
Num Resolved Nhops : 1
Num MBB Req. Nhops : 0
Num Programmed Nhops : 1
 Programmed Nhop[01] : 0.0.0.0:0, OutLabel 0 (Leaf) (StitchedFec)
Metric : 0 Mtu : 0
Num of Peers : 1
FEC Peer: 10.20.1.2:0
 Peer Flags: none (0x0)
ModTime : 01/27/16 16:39:04.097 (elapsed.: 0d 03:20:24)
->Num Egress Labels:
 None
<-Num Ingress Labels:
 <- (Label: 262139 Status: UsePop)
 Rej Status: OK
 Flow Label Tx: no, Rx: no
 Flow Label Tx Sent: no, Rx Sent: no
<Resolved as CUR Upstream>

```

## peer

**Syntax** `peer ip-address`

**Context** `tools>dump>router>ldp`

**Description** This command dumps information for an LDP peer.

## instance

**Syntax** `instance`

**Context** `tools>dump>router>ldp`

**Description** This command dumps information for the LDP instance.

## interface

**Syntax** `interface ip-int-name`

**Context** `tools>dump>router>ldp`

**Description** This command dumps information for an LDP interface.

**Parameters** *ip-int-name* — Specifies the name of an existing router.



## memory-usage

|                    |                                                      |
|--------------------|------------------------------------------------------|
| <b>Syntax</b>      | <b>memory-usage</b>                                  |
| <b>Context</b>     | tools>dump>router>ldp                                |
| <b>Description</b> | This command dumps memory usage information for LDP. |

## peer

|                    |                                                 |
|--------------------|-------------------------------------------------|
| <b>Syntax</b>      | <b>peer</b> <i>ip-address</i>                   |
| <b>Context</b>     | tools>dump>router>ldp                           |
| <b>Description</b> | This command dumps information for an LDP peer. |

## session

|                    |                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>session</b> <i>ip-addr[label-space]</i> [ <b>connection</b>   <b>peer</b>   <b>adjacency</b> ]                |
| <b>Context</b>     | tools>dump>router>ldp                                                                                            |
| <b>Description</b> | This command dumps information for an LDP session.                                                               |
| <b>Parameters</b>  | <i>ip-addr[label-space]</i> — Dumps information for the specified IP address and label space identifier.         |
| <b>Values</b>      | <i>ip-addr[label-spa*]</i> : ipv4-address:label-space<br>ipv6-address[label-space]<br>label-space - [0 to 65535] |
|                    | <b>connection</b> — Filters output for connection information.                                                   |
|                    | <b>peer</b> — Filters output for peering information.                                                            |
|                    | <b>adjacency</b> — Filters output for adjacency information.                                                     |

## sockets

|                    |                                                     |
|--------------------|-----------------------------------------------------|
| <b>Syntax</b>      | <b>sockets</b>                                      |
| <b>Context</b>     | tools>dump>router>ldp                               |
| <b>Description</b> | This command dumps information for all LDP sockets. |

# timers

|                    |                                                               |
|--------------------|---------------------------------------------------------------|
| <b>Syntax</b>      | <b>timers</b> [session <i>ip-addr</i> [ <i>label-space</i> ]] |
| <b>Context</b>     | tools>dump>router>ldp                                         |
| <b>Description</b> | This command dumps information for LDP timers.                |

# static-route

|                    |                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>static-route ldp-sync-status</b>                                                           |
| <b>Context</b>     | tools>dump>router                                                                             |
| <b>Description</b> | This command dumps the sync status of LDP interfaces that static-route tracks.                |
| <b>Parameters</b>  | <b>ldp-sync-status</b> — Displays the sync status of LDP interfaces that static-route tracks. |
| <b>Output</b>      |                                                                                               |

## Sample Output

```
*A:Dut-A# tools dump router static-route ldp-sync-status
=====
Sync Status of LDP interfaces
=====
If If Name Timer Running? Timeout Time
Index Yes/No used Left

3 ip-10.10.1.1 No 0 0
4 ip-10.10.2.1 No 0 0
=====
*A:Dut-A#
```

# ldp-sync-exit

|                    |                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ldp-sync-exit</b>                                                                       |
| <b>Context</b>     | tools>perform>router>isis                                                                  |
| <b>Description</b> | This command terminates LDP synchronization and restores actual cost of an ISIS interface. |

# run-manual-spf

|                |                           |
|----------------|---------------------------|
| <b>Syntax</b>  | <b>run-manual-spf</b>     |
| <b>Context</b> | tools>perform>router>isis |

**Description** This command runs the Shortest Path First (SPF) algorithm.

## ldp-sync-exit

**Syntax** **ldp-sync-exit**

**Context** tools>perform>router>ospf  
tools>perform>router>ospf3

**Description** This command terminates LDP synchronization and restore actual cost of an OSPF interface.

## refresh-lsas

**Syntax** **refresh-lsas** [*lsa-type*] [*area-id*]

**Context** tools>perform>router>ospf  
tools>perform>router>ospf3

**Description** This command refreshes LSAs for OSPF.

## run-manual-spf

**Syntax** **run-manual-spf** [**externals-only**]

**Context** tools>perform>router>ospf  
tools>perform>router>ospf3

**Description** This command runs the Shorted Path First (SPF) algorithm.

**Parameters** **externals-only** — Runs external only SPF.



## 6 Standards and Protocol Support



**Note:** The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

### Access Node Control Protocol (ANCP)

draft-ietf-ancp-protocol-02, *Protocol for Access Node Control Mechanism in Broadband Networks*

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

### Application Assurance (AA)

3GPP Release 12 (ADC rules over Gx interfaces)

RFC 3507, *Internet Content Adaptation Protocol (ICAP)*

### Asynchronous Transfer Mode (ATM)

AF-ILMI-0065.000, *Integrated Local Management Interface (ILMI) Version 4.0*

AF-PHY-0086.001, *Inverse Multiplexing for ATM (IMA) Specification Version 1.1*

AF-TM-0121.000, *Traffic Management Specification Version 4.1*

AF-TM-0150.00, *Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR*

GR-1113-CORE, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1*

GR-1248-CORE, *Generic Requirements for Operations of ATM Network Elements (NEs), Issue 3*

ITU-T I.432.1, *B-ISDN user-network interface - Physical layer specification: General characteristics (02/99)*

ITU-T I.610, *B-ISDN operation and maintenance principles and functions (11/95)*

RFC 1626, *Default IP MTU for use over ATM AAL5*

RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*

### Border Gateway Protocol (BGP)

draft-hares-idr-update-attr-low-bits-fix-01, *Update Attribute Flag Low Bits Clarification*

draft-ietf-idr-add-paths-04, *Advertisement of Multiple Paths in BGP*

---

draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*  
draft-ietf-idr-bgp-gr-notification-01, *Notification Message support for BGP Graceful Restart*  
draft-ietf-idr-error-handling-03, *Revised Error Handling for BGP UPDATE Messages*  
draft-ietf-sidr-origin-validation-signaling-04, *BGP Prefix Origin Validation State Extended Community*  
RFC 1772, *Application of the Border Gateway Protocol in the Internet*  
RFC 1997, *BGP Communities Attribute*  
RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*  
RFC 2439, *BGP Route Flap Damping*  
RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*  
RFC 2858, *Multiprotocol Extensions for BGP-4*  
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, *A Border Gateway Protocol 4 (BGP-4)*  
RFC 4360, *BGP Extended Communities Attribute*  
RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*  
RFC 4456, *BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)*  
RFC 4486, *Subcodes for BGP Cease Notification Message*  
RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*  
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 (helper mode)*  
RFC 4760, *Multiprotocol Extensions for BGP-4*  
RFC 4798, *Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge  
Routers (6PE)*  
RFC 4893, *BGP Support for Four-octet AS Number Space*  
RFC 5004, *Avoid BGP Best Path Transitions from One External to Another*  
RFC 5065, *Autonomous System Confederations for BGP*  
RFC 5291, *Outbound Route Filtering Capability for BGP-4*  
RFC 5575, *Dissemination of Flow Specification Rules*  
RFC 5668, *4-Octet AS Specific BGP Extended Community*  
RFC 6810, *The Resource Public Key Infrastructure (RPKI) to Router Protocol*  
RFC 6811, *Prefix Origin Validation*  
RFC 7607, *Codification of AS 0 Processing*

## Circuit Emulation

MEF-8, *Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks*, October 2004

RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*

RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)*

RFC 5287, *Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks*

## Ethernet

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*

IEEE 802.1ad, *Provider Bridges*

IEEE 802.1ag, *Connectivity Fault Management*

IEEE 802.1ah, *Provider Backbone Bridges*

IEEE 802.1ak, *Multiple Registration Protocol*

IEEE 802.1aq, *Shortest Path Bridging*

IEEE 802.1ax, *Link Aggregation*

IEEE 802.1D, *MAC Bridges*

IEEE 802.1p, *Traffic Class Expediting*

IEEE 802.1Q, *Virtual LANs*

IEEE 802.1s, *Multiple Spanning Trees*

IEEE 802.1w, *Rapid Reconfiguration of Spanning Tree*

IEEE 802.1X, *Port Based Network Access Control*

IEEE 802.3ab, *1000BASE-T*

IEEE 802.3ac, *VLAN Tag*

IEEE 802.3ad, *Link Aggregation*

IEEE 802.3ae, *10 Gb/s Ethernet*

IEEE 802.3ah, *Ethernet in the First Mile*

IEEE 802.3ba, *40 Gb/s and 100 Gb/s Ethernet*

IEEE 802.3i, *Ethernet*

IEEE 802.3u, *Fast Ethernet*

IEEE 802.3x, *Ethernet Flow Control*

IEEE 802.3z, *Gigabit Ethernet*

ITU-T G.8031/Y.1342, *Ethernet Linear Protection Switching*

ITU-T G.8032/Y.1344, *Ethernet Ring Protection Switching*

ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

---

## Ethernet VPN (EVPN)

draft-ietf-bess-evpn-overlay-02, *A Network Virtualization Overlay Solution using EVPN*

draft-ietf-bess-evpn-prefix-advertisement-02, *IP Prefix Advertisement in EVPN*

draft-ietf-bess-evpn-proxy-arp-nd-00, *Operational Aspects of Proxy-ARP/ND in EVPN Networks*

draft-ietf-bess-evpn-vpls-seamless-integ-00, *(PBB-)EVPN Seamless Integration with (PBB-)VPLS*

draft-ietf-bess-evpn-vpws-06, *VPWS support in EVPN*

RFC 7432, *BGP MPLS-Based Ethernet VPN*

RFC 7623, *Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)*

## Frame Relay

ANSI T1.617 Annex D, *DSS1 - Signalling Specification For Frame Relay Bearer Service*

FRF.1.2, *PVC User-to-Network Interface (UNI) Implementation Agreement*

FRF.12, *Frame Relay Fragmentation Implementation Agreement*

FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*

FRF.5, *Frame Relay/ATM PVC Network Interworking Implementation*

FRF2.2, *PVC Network-to-Network Interface (NNI) Implementation Agreement*

ITU-T Q.933 Annex A, *Additional procedures for Permanent Virtual Connection (PVC) status management*

## Generalized Multiprotocol Label Switching (GMPLS)

draft-ietf-ccamp-rsvp-te-srlg-collect-04, *RSVP-TE Extensions for Collecting SRLG Information*

RFC 3471, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description*

RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions*

RFC 4204, *Link Management Protocol (LMP)*

RFC 4208, *Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model*

RFC 4872, *RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery*



## **Intermediate System to Intermediate System (IS-IS)**

draft-ginsberg-isis-mi-bis-01, *IS-IS Multi-Instance* (single topology)  
draft-ietf-isis-mi-02, *IS-IS Multi-Instance*  
draft-ietf-isis-segment-routing-extensions-04, *IS-IS Extensions for Segment Routing*  
draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*  
ISO/IEC 10589:2002, Second Edition, Nov. 2002, *Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)*  
RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*  
RFC 2973, *IS-IS Mesh Groups*  
RFC 3359, *Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System*  
RFC 3719, *Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)*  
RFC 3787, *Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)*  
RFC 4971, *Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information*  
RFC 5120, *M-ISIS: Multi Topology (MT) Routing in IS-IS*  
RFC 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*  
RFC 5301, *Dynamic Hostname Exchange Mechanism for IS-IS*  
RFC 5302, *Domain-wide Prefix Distribution with Two-Level IS-IS*  
RFC 5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*  
RFC 5304, *IS-IS Cryptographic Authentication*  
RFC 5305, *IS-IS Extensions for Traffic Engineering TE*  
RFC 5306, *Restart Signaling for IS-IS (helper mode)*  
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 6213, *IS-IS BFD-Enabled TLV*  
RFC 6232, *Purge Originator Identification TLV for IS-IS*  
RFC 6233, *IS-IS Registry Extension for Purges*  
RFC 6329, *IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging*

---

## Internet Protocol (IP) — Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*

RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*

RFC 7431, *Multicast-Only Fast Reroute*

RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*

## Internet Protocol (IP) — General

draft-grant-tacacs-02, *The TACACS+ Protocol*

draft-ietf-vrrp-unified-spec-02, *Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6*

RFC 768, *User Datagram Protocol*

RFC 793, *Transmission Control Protocol*

RFC 854, *Telnet Protocol Specifications*

RFC 951, *Bootstrap Protocol (BOOTP)*

RFC 1034, *Domain Names - Concepts and Facilities*

RFC 1035, *Domain Names - Implementation and Specification*

RFC 1350, *The TFTP Protocol (revision 2)*

RFC 1534, *Interoperation between DHCP and BOOTP*

RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*

RFC 2131, *Dynamic Host Configuration Protocol*

RFC 2132, *DHCP Options and BOOTP Vendor Extensions*

RFC 2347, *TFTP Option Extension*

RFC 2348, *TFTP Blocksize Option*

RFC 2349, *TFTP Timeout Interval and Transfer Size Options*

RFC 2428, *FTP Extensions for IPv6 and NATs*

RFC 2784, *Generic Routing Encapsulation (GRE)*

RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*

RFC 2866, *RADIUS Accounting*

RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*

RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*

RFC 3046, *DHCP Relay Agent Information Option (Option 82)*

RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*

RFC 3596, *DNS Extensions to Support IP version 6*

RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*

RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*

RFC 4251, *The Secure Shell (SSH) Protocol Architecture*  
RFC 4252, *The Secure Shell (SSH) Authentication Protocol* (publickey, password)  
RFC 4253, *The Secure Shell (SSH) Transport Layer Protocol*  
RFC 4254, *The Secure Shell (SSH) Connection Protocol*  
RFC 4632, *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*  
RFC 5082, *The Generalized TTL Security Mechanism (GTSM)*  
RFC 5656, *Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer (ECDSA)*  
RFC 5880, *Bidirectional Forwarding Detection (BFD)*  
RFC 5881, *Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop)*  
RFC 5883, *Bidirectional Forwarding Detection (BFD) for Multihop Paths*  
RFC 6398, *IP Router Alert Considerations and Usage (MLD)*  
RFC 6528, *Defending against Sequence Number Attacks*  
RFC 6929, *Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions*  
RFC 7130, *Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces*

## **Internet Protocol (IP) — Multicast**

cisco-ipmulticast/pim-autorp-spec01, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast*  
draft-dolganow-bess-mvpn-expl-track-01, *Explicit Tracking with Wild Card Routes in Multicast VPN*  
draft-ietf-idmr-traceroute-ipm-07, *A "traceroute" facility for IP Multicast*  
draft-ietf-l2vpn-vpls-pim-snooping-07, *Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)*  
RFC 1112, *Host Extensions for IP Multicasting*  
RFC 2236, *Internet Group Management Protocol, Version 2*  
RFC 2375, *IPv6 Multicast Address Assignments*  
RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*  
RFC 3306, *Unicast-Prefix-based IPv6 Multicast Addresses*  
RFC 3376, *Internet Group Management Protocol, Version 3*  
RFC 3446, *Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*  
RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*  
RFC 3618, *Multicast Source Discovery Protocol (MSDP)*

- 
- RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*
- 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) (auto-RP groups)*
- RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*
- RFC 4601, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*
- RFC 4604, *Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast*
- RFC 4607, *Source-Specific Multicast for IP*
- RFC 4608, *Source-Specific Protocol Independent Multicast in 232/8*
- RFC 4610, *Anycast-RP Using Protocol Independent Multicast (PIM)*
- RFC 4611, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios*
- RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*
- RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3) / Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction*
- RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*
- RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*
- RFC 6037, *Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs*
- RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*
- RFC 6513, *Multicast in MPLS/BGP IP VPNs*
- RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/IP VPNs*
- RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs*
- RFC 6516, *IPv6 Multicast VPN (MVPN) Support Using PIM Control Plane and Selective Provider Multicast Service Interface (S-PMSI) Join Messages*
- RFC 6625, *Wildcards in Multicast VPN Auto-Discover Routes*
- RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path*
- RFC 7246, *Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context*
- RFC 7385, *IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points*

## **Internet Protocol (IP) — Version 4**

RFC 791, *Internet Protocol*  
RFC 792, *Internet Control Message Protocol*  
RFC 826, *An Ethernet Address Resolution Protocol*  
RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*  
RFC 1812, *Requirements for IPv4 Routers*  
RFC 1918, *Address Allocation for Private Internets*  
RFC 1981, *Path MTU Discovery for IP version 6*  
RFC 2003, *IP Encapsulation within IP*  
RFC 2401, *Security Architecture for Internet Protocol*  
RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*

## **Internet Protocol (IP) — Version 6**

RFC 2460, *Internet Protocol, Version 6 (IPv6) Specification*  
RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*  
RFC 2473, *Generic Packet Tunneling in IPv6 Specification*  
RFC 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*  
RFC 3122, *Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification*  
RFC 3587, *IPv6 Global Unicast Address Format*  
RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*  
RFC 3646, *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*  
RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*  
RFC 3971, *SEcure Neighbor Discovery (SEND)*  
RFC 3972, *Cryptographically Generated Addresses (CGA)*  
RFC 4007, *IPv6 Scoped Address Architecture*  
RFC 4193, *Unique Local IPv6 Unicast Addresses*  
RFC 4291, *Internet Protocol Version 6 (IPv6) Addressing Architecture*  
RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*  
RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*  
RFC 4862, *IPv6 Stateless Address Autoconfiguration (router functions)*  
RFC 4941, *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*  
RFC 5007, *DHCPv6 Leasequery*

RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*  
RFC 5952, *A Recommendation for IPv6 Address Text Representation*  
RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*  
RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*

## **Internet Protocol Security (IPsec)**

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*  
draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*  
RFC 2401, *Security Architecture for the Internet Protocol*  
RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*  
RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*  
RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*  
RFC 2406, *IP Encapsulating Security Payload (ESP)*  
RFC 2409, *The Internet Key Exchange (IKE)*  
RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*  
RFC 2560, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*  
RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman group for Internet Key Exchange (IKE)*  
RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*  
RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*  
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 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 4301, *Security Architecture for the Internet Protocol*  
RFC 4303, *IP Encapsulating Security Payload*  
RFC 4307, *Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*  
RFC 4308, *Cryptographic Suites for IPsec*  
RFC 4434, *The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)*

RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec*  
RFC 4891, *Using IPsec to Secure IPv6-in-IPv4 Tunnels*  
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 5998, *An Extension for EAP-Only Authentication in IKEv2*  
RFC 7321, *Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)*  
RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*

## **Label Distribution Protocol (LDP)**

draft-ietf-mpls-ldp-ip-pw-capability-09, *Controlling State Advertisements Of Non-negotiated LDP Applications*  
draft-ietf-mpls-ldp-ipv6-15, *Updates to LDP for IPv6*  
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, *Upstream LSR Redundancy for Multipoint LDP Tunnels*  
draft-pdutta-mpls-multi-ldp-instance-00, *Multiple LDP Instances*  
draft-pdutta-mpls-tldp-hello-reduce-04, *Targeted LDP Hello Reduction*  
RFC 3037, *LDP Applicability*  
RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol (helper mode)*  
RFC 5036, *LDP Specification*  
RFC 5283, *LDP Extension for Inter-Area Label Switched Paths (LSPs)*  
RFC 5443, *LDP IGP Synchronization*  
RFC 5561, *LDP Capabilities*  
RFC 5919, *Signaling LDP Label Advertisement Completion*  
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 6826, *Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*

## Layer Two Tunneling Protocol (L2TP) Network Server (LNS)

draft-mammoliti-l2tp-accessline-avp-04, *Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions*  
RFC 2661, *Layer Two Tunneling Protocol "L2TP"*  
RFC 2809, *Implementation of L2TP Compulsory Tunneling via RADIUS*  
RFC 3438, *Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers: Internet Assigned Numbers Authority (IANA) Considerations Update*  
RFC 3931, *Layer Two Tunneling Protocol - Version 3 (L2TPv3)*  
RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*  
RFC 4719, *Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)*  
RFC 4951, *Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"*

## Management

draft-ietf-snmpv3-update-mib-05, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*  
draft-ietf-isis-wg-mib-06, *Management Information Base for Intermediate System to Intermediate System (IS-IS)*  
draft-ietf-mboned-msdp-mib-01, *Multicast Source Discovery protocol MIB*  
draft-ietf-mpls-ldp-mib-07, *Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP)*  
draft-ietf-mpls-lsr-mib-06, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2*  
draft-ietf-mpls-te-mib-04, *Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base*  
draft-ietf-ospf-mib-update-08, *OSPF Version 2 Management Information Base*  
ianaaddressfamilynumbers-mib, *IANA-ADDRESS-FAMILY-NUMBERS-MIB*  
ianagmplstc-mib, *IANA-GMPLS-TC-MIB*  
ianaiftype-mib, *IANAIfType-MIB*  
ianaiprouteprotocol-mib, *IANA-RTPROTO-MIB*  
IEEE8021-CFM-MIB, *IEEE P802.1ag(TM) CFM MIB*  
IEEE8021-PAE-MIB, *IEEE 802.1X MIB*  
IEEE8023-LAG-MIB, *IEEE 802.3ad MIB*  
LLDP-MIB, *IEEE P802.1AB(TM) LLDP MIB*  
RFC 1157, *A Simple Network Management Protocol (SNMP)*  
RFC 1212, *Concise MIB Definitions*



RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*

RFC 1215, *A Convention for Defining Traps for use with the SNMP*

RFC 1724, *RIP Version 2 MIB Extension*

RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIv2*

RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIv2*

RFC 2206, *RSVP Management Information Base using SMIv2*

RFC 2213, *Integrated Services Management Information Base using SMIv2*

RFC 2494, *Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type*

RFC 2514, *Definitions of Textual Conventions and OBJECT-IDENTITIES for ATM Management*

RFC 2515, *Definitions of Managed Objects for ATM Management*

RFC 2570, *SNMP Version 3 Framework*

RFC 2571, *An Architecture for Describing SNMP Management Frameworks*

RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*

RFC 2573, *SNMP Applications*

RFC 2574, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*

RFC 2575, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*

RFC 2578, *Structure of Management Information Version 2 (SMIv2)*

RFC 2579, *Textual Conventions for SMIv2*

RFC 2580, *Conformance Statements for SMIv2*

RFC 2787, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol*

RFC 2819, *Remote Network Monitoring Management Information Base*

RFC 2856, *Textual Conventions for Additional High Capacity Data Types*

RFC 2863, *The Interfaces Group MIB*

RFC 2864, *The Inverted Stack Table Extension to the Interfaces Group MIB*

RFC 2933, *Internet Group Management Protocol MIB*

RFC 3014, *Notification Log MIB*

RFC 3164, *The BSD syslog Protocol*

RFC 3165, *Definitions of Managed Objects for the Delegation of Management Scripts*

RFC 3231, *Definitions of Managed Objects for Scheduling Management Operations*

---

RFC 3273, *Remote Network Monitoring Management Information Base for High Capacity Networks*

RFC 3416, *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*

RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4)*

RFC 3419, *Textual Conventions for Transport Addresses*

RFC 3498, *Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures*

RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework*

RFC 3592, *Definitions of Managed Objects for the Synchronous Optical Network/ Synchronous Digital Hierarchy (SONET/SDH) Interface Type*

RFC 3593, *Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals*

RFC 3635, *Definitions of Managed Objects for the Ethernet-like Interface Types*

RFC 3637, *Definitions of Managed Objects for the Ethernet WAN Interface Sublayer*

RFC 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*

RFC 3877, *Alarm Management Information Base (MIB)*

RFC 3895, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*

RFC 3896, *Definitions of Managed Objects for the DS3/E3 Interface Type*

RFC 4001, *Textual Conventions for Internet Network Addresses*

RFC 4022, *Management Information Base for the Transmission Control Protocol (TCP)*

RFC 4113, *Management Information Base for the User Datagram Protocol (UDP)*

RFC 4220, *Traffic Engineering Link Management Information Base*

RFC 4273, *Definitions of Managed Objects for BGP-4*

RFC 4292, *IP Forwarding Table MIB*

RFC 4293, *Management Information Base for the Internet Protocol (IP)*

RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*

RFC 4511, *Lightweight Directory Access Protocol (LDAP): The Protocol*

RFC 4513, *Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms (TLS)*

RFC 4631, *Link Management Protocol (LMP) Management Information Base (MIB)*

RFC 4878, *Definitions and Managed Objects for Operations, Administration, and Maintenance (OAM) Functions on Ethernet-Like Interfaces*

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*  
RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2* (TLS client, RSA public key)  
RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP)* (server, unauthenticated mode)  
RFC 5938, *Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)*  
RFC 6020, *YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)*  
RFC 6038, *Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features*  
RFC 6241, *Network Configuration Protocol (NETCONF)*  
RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*  
RFC 6243, *With-defaults Capability for NETCONF*  
RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*  
RFC 6425, *Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping*  
RFC 7420, *Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module*  
SFLOW-MIB, *sFlow MIB Version 1.3 (Draft 5)*

## **Multiprotocol Label Switching - Transport Profile (MPLS-TP)**

RFC 5586, *MPLS Generic Associated Channel*  
RFC 5921, *A Framework for MPLS in Transport Networks*  
RFC 5960, *MPLS Transport Profile Data Plane Architecture*  
RFC 6370, *MPLS Transport Profile (MPLS-TP) Identifiers*  
RFC 6378, *MPLS Transport Profile (MPLS-TP) Linear Protection*  
RFC 6426, *MPLS On-Demand Connectivity and Route Tracing*  
RFC 6427, *MPLS Fault Management Operations, Administration, and Maintenance (OAM)*  
RFC 6428, *Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile*  
RFC 6478, *Pseudowire Status for Static Pseudowires*  
RFC 7213, *MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing*

---

## Multiprotocol Label Switching (MPLS)

RFC 3031, *Multiprotocol Label Switching Architecture*

RFC 3032, *MPLS Label Stack Encoding*

RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*

RFC 4023, *Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)*

RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*

RFC 5332, *MPLS Multicast Encapsulations*

RFC 6790, *The Use of Entropy Labels in MPLS Forwarding*

## Network Address Translation (NAT)

draft-ietf-behave-address-format-10, *IPv6 Addressing of IPv4/IPv6 Translators*

draft-ietf-behave-v6v4-xlate-23, *IP/ICMP Translation Algorithm*

draft-miles-behave-l2nat-00, *Layer2-Aware NAT*

draft-nishitani-cgn-02, *Common Functions of Large Scale NAT (LSN)*

RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*

RFC 5382, *NAT Behavioral Requirements for TCP*

RFC 5508, *NAT Behavioral Requirements for ICMP*

RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*

RFC 6333, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*

RFC 6334, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite*

RFC 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*

RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*

## Open Shortest Path First (OSPF)

draft-ietf-ospf-prefix-link-attr-06, *OSPFv2 Prefix/Link Attribute Advertisement*

draft-ietf-ospf-segment-routing-extensions-04, *OSPF Extensions for Segment Routing*

RFC 1586, *Guidelines for Running OSPF Over Frame Relay Networks*

RFC 1765, *OSPF Database Overflow*

RFC 2328, *OSPF Version 2*

RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*

RFC 3509, *Alternative Implementations of OSPF Area Border Routers*

RFC 3623, *Graceful OSPF Restart Graceful OSPF Restart (helper mode)*  
RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*  
RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*  
RFC 4222, *Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance*  
RFC 4552, *Authentication/Confidentiality for OSPFv3*  
RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*  
RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*  
RFC 4970, *Extensions to OSPF for Advertising Optional Router Capabilities*  
RFC 5185, *OSPF Multi-Area Adjacency*  
RFC 5187, *OSPFv3 Graceful Restart (helper mode)*  
RFC 5243, *OSPF Database Exchange Summary List Optimization*  
RFC 5250, *The OSPF Opaque LSA Option*  
RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*  
RFC 5340, *OSPF for IPv6*  
RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*  
RFC 5838, *Support of Address Families in OSPFv3*  
RFC 6987, *OSPF Stub Router Advertisement*

## **OpenFlow**

ONF *OpenFlow Switch Specification Version 1.3.1* (OpenFlow-hybrid switches)

## **Path Computation Element Protocol (PCEP)**

draft-alvarez-pce-path-profiles-04, *PCE Path Profiles*  
draft-ietf-pce-segment-routing-05, *PCEP Extensions for Segment Routing*  
draft-ietf-pce-stateful-pce-11, *PCEP Extensions for Stateful PCE*  
RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*

## **Point-to-Point Protocol (PPP)**

RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*  
RFC 1377, *The PPP OSI Network Layer Control Protocol (OSINLCP)*  
RFC 1661, *The Point-to-Point Protocol (PPP)*  
RFC 1662, *PPP in HDLC-like Framing*

RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*  
RFC 1989, *PPP Link Quality Monitoring*  
RFC 1990, *The PPP Multilink Protocol (MP)*  
RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*  
RFC 2153, *PPP Vendor Extensions*  
RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*  
RFC 2615, *PPP over SONET/SDH*  
RFC 2686, *The Multi-Class Extension to Multi-Link PPP*  
RFC 2878, *PPP Bridging Control Protocol (BCP)*  
RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*  
RFC 5072, *IP Version 6 over PPP*

## **Policy Management and Credit Control**

3GPP TS 29.212 Release 11, *Policy and Charging Control (PCC); Reference points (Gx support as it applies to wireline environment (BNG))*  
RFC 3588, *Diameter Base Protocol*  
RFC 4006, *Diameter Credit-Control Application*

## **Pseudowire**

draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS Interworking*  
MFA Forum 12.0.0, *Multiservice Interworking - Ethernet over MPLS*  
MFA Forum 13.0.0, *Fault Management for Multiservice Interworking v1.0*  
MFA Forum 16.0.0, *Multiservice Interworking - IP over MPLS*  
MFA Forum 9.0.0, *The Use of Virtual trunks for ATM/MPLS Control Plane Interworking*  
RFC 3916, *Requirements for Pseudo- Wire Emulation Edge-to-Edge (PWE3)*  
RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3)*  
RFC 4385, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*  
RFC 4446, *IANA Allocations for Pseudowire Edge to Edge Emulation (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 4619, *Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks*

- RFC 4717, *Encapsulation Methods for Transport Asynchronous Transfer Mode (ATM) over 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 5659, *An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge*
- RFC 5885, *Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)*
- RFC 6073, *Segmented Pseudowire*
- RFC 6310, *Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping*
- RFC 6391, *Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network*
- RFC 6575, *Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs*
- RFC 6718, *Pseudowire Redundancy*
- RFC 6829, *Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6*
- RFC 6870, *Pseudowire Preferential Forwarding Status bit*
- RFC 7023, *MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking*
- RFC 7267, *Dynamic Placement of Multi-Segment Pseudowires*

## **Quality of Service (QoS)**

- RFC 2430, *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*
- RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*
- RFC 2598, *An Expedited Forwarding PHB*
- RFC 3140, *Per Hop Behavior Identification Codes*
- RFC 3260, *New Terminology and Clarifications for Diffserv*

## **Resource Reservation Protocol - Traffic Engineering (RSVP-TE)**

- draft-newton-mpls-te-dynamic-overbooking-00, *A Diffserv-TE Implementation Model to dynamically change booking factors during failure events*
- 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, *RSVP-TE: Extensions to RSVP for LSP Tunnels*  
RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling  
Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions  
(IF\_ID RSVP\_HOP object with unnumbered interfaces and RSVP-TE graceful  
restart helper procedures)*  
RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic  
Engineering (RSVP-TE)*  
RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic  
Engineering*  
RFC 3906, *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic  
Engineering Tunnels*  
RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*  
RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic  
Engineering*  
RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware  
MPLS Traffic Engineering*  
RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS  
Traffic Engineering*  
RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*  
RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering  
(RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*  
RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*  
RFC 5151, *Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource  
Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*  
RFC 5712, *MPLS Traffic Engineering Soft Preemption*  
RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering  
Networks*

## **Routing Information Protocol (RIP)**

RFC 1058, *Routing Information Protocol*  
RFC 2080, *RIPng for IPv6*  
RFC 2082, *RIP-2 MD5 Authentication*  
RFC 2453, *RIP Version 2*



## **Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH)**

ITU-G.841, *Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1, issued in July 2002*

### **Timing**

GR-1244-CORE, *Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005*

GR-253-CORE, *SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000*

IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*

ITU-T G.781, *Synchronization layer functions, issued 09/2008*

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003*

ITU-T G.8261, *Timing and synchronization aspects in packet networks, issued 04/2008*

ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007*

ITU-T G.8264, *Distribution of timing information through packet networks, issued 10/2008*

ITU-T G.8265.1, *Precision time protocol telecom profile for frequency synchronization, issued 10/2010*

ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014*

RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*

### **Virtual Private LAN Service (VPLS)**

RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*

RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*

RFC 5501, *Requirements for Multicast Support in Virtual Private LAN Services*

RFC 6074, *Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)*

RFC 7041, *Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging*

---

RFC 7117, *Multicast in Virtual Private LAN Service (VPLS)*

## **Voice and Video**

DVB BlueBook A86, *Transport of MPEG-2 TS Based DVB Services over IP Based Networks*

ETSI TS 101 329-5 Annex E, *QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring*

ITU-T G.1020 Appendix I, *Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks - Mean Absolute Packet Delay Variation & Markov Models*

ITU-T G.107, *The E Model - A computational model for use in planning*

ITU-T P.564, *Conformance testing for voice over IP transmission quality assessment models*

RFC 3550 Appendix A.8, *RTP: A Transport Protocol for Real-Time Applications (estimating the interarrival jitter)*

RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*

RFC 4588, *RTP Retransmission Payload Format*

## **Wireless Local Area Network (WLAN) Gateway**

3GPP TS 23.402, *Architecture enhancements for non-3GPP accesses (S2a roaming based on GPRS)*

# Customer Document and Product Support



## Customer Documentation

[Customer Documentation Welcome Page](#)



## Technical Support

[Product Support Portal](#)



## Documentation Feedback

[Customer Documentation Feedback](#)

