



Junos[®] OS

MPLS Applications Feature Guide for Routing Devices

Release
15.1



Modified: 2015-10-08

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Junos[®] OS MPLS Applications Feature Guide for Routing Devices

15.1

Copyright © 2015, Juniper Networks, Inc.
All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About the Documentation	xxxvii
	Documentation and Release Notes	xxxvii
	Supported Platforms	xxxvii
	Using the Examples in This Manual	xxxvii
	Merging a Full Example	xxxviii
	Merging a Snippet	xxxviii
	Documentation Conventions	xxxix
	Documentation Feedback	xli
	Requesting Technical Support	xli
	Self-Help Online Tools and Resources	xli
	Opening a Case with JTAC	xlii
Part 1	Understanding Traffic Engineering	
Chapter 1	Traffic Engineering Overview	3
	Traffic Engineering Capabilities	3
	Components of Traffic Engineering	4
	Packet Forwarding Component	4
	Packet Forwarding Based on Label Swapping	4
	How a Packet Traverses an MPLS Backbone	5
	Information Distribution Component	5
	Path Selection Component	6
	Offline Path Planning and Analysis	6
	Signaling Component	7
	Flexible LSP Calculation and Configuration	7
	Link-State Distribution Using BGP Overview	8
	Role of an Interior Gateway Protocol	8
	Limitations of an Interior Gateway Protocol	9
	Need for Spanning Link-State Distribution	9
	Using BGP as a Solution	10
	Overview	10
	Implementation	11
	Supported and Unsupported Features	15
Part 2	Configuring MPLS	
Chapter 2	MPLS Overview	19
	Introduction to MPLS	20
	Supported MPLS Standards	20
	Link-Layer Support in MPLS	23
	MPLS and Traffic Engineering	24

MPLS Label Overview	24
Special MPLS Labels	25
MPLS Label Allocation	26
Operations on MPLS Labels	27
Entropy Label Support in Mixed Mode Overview	28
Routers in an LSP	28
How a Packet Travels Along an LSP	28
Types of LSPs	29
Scope of LSPs	29
Constrained-Path LSP Computation	29
How CSPF Selects a Path	31
CSPF Path Selection Tie-Breaking	32
Computing CSPF Paths Offline	33
Path Computation for LSPs on an Overloaded Router	33
Computing Backup Paths for LSPs Using Fate Sharing	34
Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts . . .	34
Enabling IGP Shortcuts	36
LSPs Qualified in IGP Shortcut Computations	36
IGP Shortcut Applications	36
IGP Shortcuts and Routing Tables	37
IGP Shortcuts and VPNs	38
Advertising LSPs into IGP	38
IP and MPLS Packets on Aggregated Interfaces	39
MPLS Applications	40
BGP Destinations	40
IGP and BGP Destinations	41
Selecting a Forwarding LSP Next Hop	42
Example: Assigning Different Forwarding Next-Hop LSPs to Different Destination Prefixes	42
MPLS and Routing Tables	43
MPLS and Traffic Protection	45
Fast Reroute Overview	46
Detour Merging Process	48
Detour Computations	49
Fast Reroute Path Optimization	50
On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview . . .	50
Importance of Measuring Packet Loss and Delay	50
Defining Packet Loss, Delay, and Throughput	51
Packet Loss and Delay Measurement Mechanisms	51
Packet Loss and Delay Metrics	52
Packet Loss and Delay Measurement Concepts	52
Packet Loss and Delay Measurement Functionality	55
Packet Loss and Delay Features	56

Chapter 3	Configuring MPLS Routers	59
	Minimum MPLS Configuration	59
	Configuring the Ingress Router for MPLS-Signaled LSPs	60
	Creating Named Paths	60
	Examples: Creating Named Paths	62
	Configuring Alternate Backup Paths Using Fate Sharing	62
	Configuring Fate Sharing	63
	Implications for CSPF	64
	Implications for CSPF When Fate Sharing with Bypass LSPs	64
	Example: Configuring Fate Sharing	64
	Example: Configuring an Explicit-Path LSP	65
	Example: Configuring a Constrained-Path LSP for Which Junos OS Makes All	
	Forwarding Decisions	66
	Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most	
	Forwarding Decisions and Considers Hop Constraints	66
	Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most	
	Forwarding Decisions and the Secondary Path Is Explicit	67
	Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs	68
	Improving Traffic Engineering Database Accuracy with RSVP PathErr	
	Messages	68
	PathErr Messages	68
	Identifying the Problem Link	69
	Configuring the Router to Improve Traffic Engineering Database	
	Accuracy	69
	Configuring MPLS-Signaled LSPs to Use GRE Tunnels	70
	Example: Configuring MPLS-Signaled LSPs to Use GRE Tunnels	70
	Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks	71
	SRLG Overview	80
	Example: Configuring SRLG	81
	Example: Excluding SRLG Links Completely for the Secondary LSP	90
	Example: Configuring SRLG with Link Protection	95
	Example: Configuring SRLG with Link Protection with the exclude-srlg	
	Option	116
	Configuring the MPLS Transport Profile for OAM	136
	MPLS Transport Profile Overview	136
	Example: Configuring the MPLS Transport Profile for OAM	136
	Configuring MPLS Inter-AS Link-Node Protection with Labeled BGP	148
	Understanding MPLS Inter-AS Link Protection	148
	Example: Configuring MPLS Inter-AS Link-Node Protection	150
	Configuring Egress Protection Service Mirroring for BGP Signaled Layer 2	
	Services	164
	Example: Configuring MPLS Egress Protection Service Mirroring for BGP Signaled	
	Layer 2 Services	168
	Example: Configuring Layer 3 VPN Egress Protection with PLR as Protector	184
Chapter 4	Configuring MPLS-Signaled LSPs	211
	Configuring the Ingress and Egress Router Addresses for LSPs	212
	Configuring the Ingress Router Address for LSPs	212
	Configuring the Egress Router Address for LSPs	212

Preventing the Addition of Egress Router Addresses to Routing Tables	213
Configuring Primary and Secondary LSPs	214
Configuring Primary and Secondary Paths for an LSP	214
Configuring the Revert Timer for LSPs	215
Specifying the Conditions for Path Selection	216
Configuring a Text Description for LSPs	217
Configuring the Entropy Label for LSPs	218
Configuring Corouted Bidirectional LSPs	220
Configuring Ultimate-Hop Popping for LSPs	222
Configuring an LSP Across ASs	225
Configuring Fast Reroute	226
Configuring the Optimization Interval for Fast Reroute Paths	228
Adding LSP-Related Routes to the inet.3 or inet6.3 Routing Table	228
Configuring the Connection Between Ingress and Egress Routers	229
Configuring LSP Metrics	230
Configuring Dynamic LSP Metrics	230
Configuring Static LSP Metrics	231
Configuring CSPF Tie Breaking	232
Configuring Load Balancing Based on MPLS Labels on DPC I-Chip-Based Hardware	232
Disabling Normal TTL Decrementing	236
Configuring MPLS Soft Preemption	238
Disabling Constrained-Path LSP Computation	239
Configuring Administrative Groups for LSPs	240
Configuring Extended Administrative Groups for LSPs	242
Configuring Preference Values for LSPs	243
Disabling Path Route Recording by LSPs	244
Configuring Class of Service for MPLS LSPs	244
Class of Service for MPLS Overview	244
Configuring the MPLS CoS Bits	245
Rewriting IEEE 802.1p Packet Headers with the MPLS CoS Value	246
Achieving a Make-Before-Break, Hitless Switchover for LSPs	246
Specifying the Amount of Time the Router Waits to Switch Over to New Paths	247
Specifying the Amount of Time to Delay the Tear Down of Old Paths	248
Achieving a Hitless, MBB Switchover Without Artificial Delays	248
Configuring Adaptive LSPs	249
Configuring Priority and Preemption for LSPs	250
Optimizing Signaled LSPs	251
Configuring the Smart Optimize Timer for LSPs	255
Limiting the Number of Hops in LSPs	256
Configuring the Bandwidth Value for LSPs	256
Automatic Bandwidth Allocation for LSPs	257
Configuring Automatic Bandwidth Allocation for LSPs	257
Configuring Automatic Bandwidth Allocation on LSPs	258
Configuring the Automatic Bandwidth Allocation Interval	259
Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth	259
Configuring the Automatic Bandwidth Adjustment Threshold	260

	Configuring a Limit on Bandwidth Overflow and Underflow	
	Samples	261
	Configuring Passive Bandwidth Utilization Monitoring	263
	Requesting Automatic Bandwidth Allocation Adjustment	263
	Configuring Reporting of Automatic Bandwidth Allocation Statistics for	
	LSPs	264
	Configuring Hot Standby of Secondary Paths for LSPs	267
	Damping Advertisement of LSP State Changes	268
Chapter 5	Configuring Static and Explicit-Path LSPs	271
	Configuring Static LSPs	271
	Configuring the Ingress Router for Static LSPs	271
	Example: Configuring the Ingress Router	273
	Configuring the Intermediate (Transit) and Egress Routers for Static	
	LSPs	274
	Example: Configuring an Intermediate Router	275
	Example: Configuring an Egress Router	276
	Configuring a Bypass LSP for the Static LSP	276
	Configuring the Protection Revert Timer for Static LSPs	277
	Configuring Static Unicast Routes for Point-to-Multipoint LSPs	277
	Configuring Explicit-Path LSPs	278
Chapter 6	Configuring Point-to-Multipoint LSPs	281
	Point-to-Multipoint LSPs Overview	281
	Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs	283
	Configuring the Primary Point-to-Multipoint LSP	283
	Configuring a Branch LSP for Point-to-Multipoint LSPs	283
	Configuring the Branch LSP as a Dynamic Path	284
	Configuring the Branch LSP as a Static Path	284
	Example: Configuring a Collection of Paths to Create an RSVP-Signaled	
	Point-to-Multipoint LSP	285
	Configuring Inter-Domain Point-to-Multipoint LSPs	303
	Configuring Link Protection for Point-to-Multipoint LSPs	304
	Configuring Graceful Restart for Point-to-Multipoint LSPs	305
	Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs	306
	Example: Configuring Multicast RPF Check Policy for a Point-to-Multipoint	
	LSP	306
	Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs	307
	Enabling Point-to-Point LSPs to Monitor Egress PE Routers	307
	Preserving Point-to-Multipoint LSP Functioning with Different Junos OS	
	Releases	308
Chapter 7	Configuring DiffServ-Aware Traffic Engineering to Achieve Service Level	
	Guarantees on an MPLS network	309
	DiffServ-Aware Traffic Engineering Introduction	310
	DiffServ-Aware Traffic Engineering Standards	310
	DiffServ-Aware Traffic Engineering Terminology	310
	DiffServ-Aware Traffic Engineering Features	311
	DiffServ-Aware Traffic Engineered LSPs	312
	DiffServ-Aware Traffic Engineered LSPs Overview	312

DiffServ-Aware Traffic Engineered LSPs Operation	313
Multiclass LSPs	313
Multiclass LSP Overview	314
Establishing a Multiclass LSP on the Differentiated Services Domain	314
Configuring Routers for DiffServ-Aware Traffic Engineering	315
Configuring the Bandwidth Model	316
Configuring Traffic Engineering Classes	317
Requirements and Limitations for the Traffic Engineering Class	
Matrix	318
Configuring Class of Service for DiffServ-Aware Traffic Engineering	318
LSP Bandwidth Oversubscription Overview	319
LSP Size Oversubscription	320
LSP Link Size Oversubscription	320
Class Type Oversubscription and Local Oversubscription Multipliers	320
Class Type Bandwidth and the LOM	321
LOM Calculation for the MAM and Extended MAM Bandwidth Models	321
LOM Calculation for the Russian Dolls Bandwidth Model	322
Example: LOM Calculation	322
Configuring the Bandwidth Subscription Percentage for LSPs	323
Constraints on Configuring Bandwidth Subscription	324
Configuring LSPs for DiffServ-Aware Traffic Engineering	325
Configuring Class of Service for the Interfaces	326
Configuring IGP	326
Configuring Traffic-Engineered LSPs	326
Configuring Policing for LSPs	327
Configuring Fast Reroute for Traffic-Engineered LSPs	327
Configuring Multiclass LSPs	328
Configuring Class of Service for the Interfaces	328
Configuring the IGP	329
Configuring Class-Type Bandwidth Constraints for Multiclass LSPs	329
Configuring Policing for Multiclass LSPs	330
Configuring Fast Reroute for Multiclass LSPs	330
Chapter 8	
Configuring Miscellaneous MPLS Properties	333
Configuring the Maximum Number of MPLS Labels	334
Configuring MPLS to Pop the Label on the Ultimate-Hop Router	335
Advertising Explicit Null Labels to BGP Peers	336
Configuring Traffic Engineering for LSPs	336
Using LSPs for Both BGP and IGP Traffic Forwarding	337
Using LSPs for Forwarding in Virtual Private Networks	337
Using RSVP and LDP Routes for Forwarding but Not Route Selection	338
Advertising the LSP Metric in Summary LSAs	339
Enabling Interarea Traffic Engineering	339
Enabling Inter-AS Traffic Engineering for LSPs	340
Inter-AS Traffic Engineering Requirements	340
Inter-AS Traffic Engineering Limitations	341
Configuring OSPF Passive TE Mode	342
Configuring MPLS to Gather Statistics	342
Configuring System Log Messages and SNMP Traps for LSPs	344

Configuring MPLS Firewall Filters and Policers	345
Configuring MPLS Firewall Filters	345
Examples: Configuring MPLS Firewall Filters	346
Configuring Policers for LSPs	347
LSP Policer Limitations	348
Example: Configuring an LSP Policer	348
Configuring Automatic Policers	349
Configuring Automatic Policers for LSPs	350
Configuring Automatic Policers for DiffServ-Aware Traffic Engineering LSPs	351
Configuring Automatic Policers for Point-to-Multipoint LSPs	351
Disabling Automatic Policing on an LSP	352
Example: Configuring Automatic Policing for an LSP	352
Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets	352
Configuring MPLS Rewrite Rules	353
Rewriting the EXP Bits of All Three Labels of an Outgoing Packet	353
Rewriting MPLS and IPv4 Packet Headers	353
Configuring BFD for MPLS IPv4 LSPs	354
Configuring BFD for RSVP-Signaled LSPs	355
Configuring a Failure Action for the BFD Session on an RSVP LSP	356
BFD-Triggered Local Repair for Rapid Convergence	357
Understanding BFD-Triggered Local Protection	357
Purpose of BFD-Triggered Local Repair	357
Configuring BFD-Triggered Local Repair	358
Disabling BFD-Triggered Local Repair	358
Disabling BFD-Triggered Local Repair	358
Pinging LSPs	359
Pinging MPLS LSPs	359
Pinging Point-to-Multipoint LSPs	360
Pinging the Endpoint Address of MPLS LSPs	360
Pinging CCC LSPs	360
Pinging Layer 3 VPNs	360
Support for LSP Ping and Traceroute Commands Based on RFC 4379	360
Tracing MPLS and LSP Packets and Operations	361
Configuring Link State Distribution Using BGP	362
Example: Configuring Link State Distribution Using BGP	364
Dynamic Bandwidth Management Using Container LSP Overview	382
Understanding RSVP Multipath Extensions	382
Junos OS RSVP Multipath Implementation	383
Current Traffic Engineering Challenges	383
Using Container LSP as a Solution	386
Accommodating the New Demand X	387
Creating New LSPs to Meet Demand X	387
Assigning Bandwidth to the New LSPs	387
Controlling the LSP Paths	387
Junos OS Container LSP Implementation	388
Container LSP Terminology	388
LSP Splitting	389
LSP Merging	391

	Node and Link Protection	393
	Naming Convention	393
	Normalization	394
	Constraint-Based Routing Path Computation	399
	Sampling	400
	Support for NSR, IPG-FA, and Static Routes	400
	Configuration Statements Supported for Container LSPs	403
	Impact of Configuring Container LSPs on Network Performance	407
	Supported and Unsupported Features	408
	Configuring Dynamic Bandwidth Management Using Container LSP	409
	Example: Configuring Dynamic Bandwidth Management Using Container LSP	413
	Configuring On-Demand Loss and Delay Measurement	438
	Example: Configuring On-Demand Loss and Delay Measurement	439
	Configuring Pro-Active Loss and Delay Measurements	448
	Example: Configuring Pro-active Loss and Delay Measurements for Bidirectional MPLS LSPs	450
Part 3	Configuring RSVP	
Chapter 9	RSVP Overview	459
	RSVP Overview	459
	Supported RSVP Standards	460
	Junos OS RSVP Protocol Implementation	461
	RSVP Operation Overview	462
	RSVP Authentication	462
	RSVP and IGP Hello Packets and Timers	462
	RSVP Message Types	463
	Understanding RSVP Automatic Mesh	463
	Path Messages	465
	Resv Messages	465
	PathTear Messages	465
	ResvTear Messages	465
	PathErr Messages	466
	ResvErr Messages	466
	ResvConfirm Messages	466
	RSVP Reservation Styles	466
	RSVP Refresh Reduction	467
	MTU Signaling in RSVP	468
	How the Correct MTU Is Signaled in RSVP	469
	Determining an Outgoing MTU Value	469
	MTU Signaling in RSVP Limitations	470
Chapter 10	Configuring RSVP	471
	Minimum RSVP Configuration	471
	Configuring RSVP and MPLS	472
	Example: Configuring RSVP and MPLS	472

	Configuring RSVP Interfaces	473
	Configuring RSVP Refresh Reduction	473
	Determining the Refresh Reduction Capability of RSVP Neighbors	475
	Configuring the RSVP Hello Interval	475
	Configuring RSVP Authentication	476
	Configuring the Bandwidth Subscription for Class Types	476
	Configuring the RSVP Update Threshold on an Interface	476
	Configuring RSVP for Unnumbered Interfaces	477
	Configuring RSVP Node ID Hellos	478
	Configuring Hello Acknowledgments for Nonsession RSVP Neighbors	479
	Switching LSPs Away from a Network Node	479
	Configuring RSVP Setup Protection	480
	Configuring Load Balancing Across RSVP LSPs	481
	Configuring RSVP Automatic Mesh	482
	Configuring Timers for RSVP Refresh Messages	483
	Preempting RSVP Sessions	484
	Configuring MTU Signaling in RSVP	484
	Enabling MTU Signaling in RSVP	485
	Enabling Packet Fragmentation	485
	Configuring Ultimate-Hop Popping for LSPs	486
	Configuring RSVP to Pop the Label on the Ultimate-Hop Router	489
	Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs	490
	Tracing RSVP Protocol Traffic	491
	Examples: Tracing RSVP Protocol Traffic	492
Chapter 11	Configuring RSVP Link Protection and Node Protection to Protect from Traffic Failures	495
	Link Protection	495
	Multiple Bypass LSPs for Link Protection	496
	Node Protection	497
	Fast Reroute, Node Protection, and Link Protection	498
	LSP Protection Overview	498
	LSP Protection Types Comparison	499
	One-to-One Backup Implementation	499
	Facility Backup Implementation	500
	Configuring Link Protection on Interfaces Used by LSPs	502
	Configuring Bypass LSPs	503
	Configuring the Next-Hop or Next-Next-Hop Node Address for Bypass LSPs	504
	Configuring Administrative Groups for Bypass LSPs	504
	Configuring the Bandwidth for Bypass LSPs	504
	Configuring Class of Service for Bypass LSPs	505
	Configuring the Hop Limit for Bypass LSPs	505
	Configuring the Maximum Number of Bypass LSPs	506
	Disabling CSPF for Bypass LSPs	506
	Disabling Node Protection for Bypass LSPs	507
	Configuring the Optimization Interval for Bypass LSPs	507
	Configuring an Explicit Path for Bypass LSPs	508
	Configuring the Amount of Bandwidth Subscribed for Bypass LSPs	508

	Configuring Priority and Preemption for Bypass LSPs	509
	Configuring Node Protection or Link Protection for LSPs	509
	Configuring Inter-AS Node and Link Protection	510
Chapter 12	Configuring RSVP Graceful Restart for High Availability	511
	RSVP Graceful Restart	511
	RSVP Graceful Restart Standard	511
	RSVP Graceful Restart Terminology	512
	RSVP Graceful Restart Operation	512
	Processing the Restart Cap Object	513
	Configuring RSVP Graceful Restart	514
	Enabling Graceful Restart for All Routing Protocols	514
	Disabling Graceful Restart for RSVP	514
	Disabling RSVP Helper Mode	515
	Configuring the Maximum Helper Recovery Time	515
	Configuring the Maximum Helper Restart Time	515
Part 4	Configuring LDP	
Chapter 13	LDP Overview	519
	LDP Introduction	519
	Supported LDP Standards	520
	Junos OS LDP Protocol Implementation	520
	LDP Operation	521
	LDP Message Types	521
	Discovery Messages	521
	Session Messages	522
	Advertisement Messages	522
	Notification Messages	522
	Tunneling LDP LSPs in RSVP LSPs	523
	Tunneling LDP LSPs in RSVP LSPs Overview	523
	Label Operations	523
	LDP Session Protection	525
Chapter 14	Configuring LDP	527
	Minimum LDP Configuration	528
	Enabling and Disabling LDP	528
	Configuring the LDP Timer for Hello Messages	528
	Configuring the LDP Timer for Link Hello Messages	529
	Configuring the LDP Timer for Targeted Hello Messages	529
	Configuring the Delay Before LDP Neighbors Are Considered Down	529
	Configuring the LDP Hold Time for Link Hello Messages	530
	Configuring the LDP Hold Time for Targeted Hello Messages	530
	Enabling Strict Targeted Hello Messages for LDP	531
	Configuring the Interval for LDP Keepalive Messages	531
	Configuring the LDP Keepalive Timeout	531
	Configuring LDP Route Preferences	532
	LDP Graceful Restart	532

Configuring LDP Graceful Restart	533
Enabling Graceful Restart	533
Disabling LDP Graceful Restart or Helper Mode	534
Configuring Reconnect Time	534
Configuring Recovery Time and Maximum Recovery Time	535
Filtering Inbound LDP Label Bindings	535
Examples: Filtering Inbound LDP Label Bindings	536
Filtering Outbound LDP Label Bindings	537
Examples: Filtering Outbound LDP Label Bindings	538
Specifying the Transport Address Used by LDP	539
Configuring the Prefixes Advertised into LDP from the Routing Table	540
Example: Configuring the Prefixes Advertised into LDP	540
Configuring FEC Deaggregation	541
Configuring Policers for LDP FECs	541
Configuring LDP IPv4 FEC Filtering	542
Configuring BFD for LDP LSPs	543
Configuring ECMP-Aware BFD for LDP LSPs	546
Configuring a Failure Action for the BFD Session on an LDP LSP	546
Configuring the Holddown Interval for the BFD Session	547
Configuring OAM Ingress Policies for LDP	547
Configuring LDP Link Protection	548
Example: Configuring LDP Link Protection	549
LDP Link Protection Overview	549
Introduction to LDP	550
Junos OS LDP Protocol Implementation	550
Understanding Multipoint Extensions to LDP	550
Using Multipoint Extensions to LDP on Targeted LDP Sessions	551
Current Limitations of LDP Link Protection	552
Using RSVP LSP as a Solution	553
Understanding Multicast LDP Link Protection	555
Different Modes for Providing LDP Link Protection	555
Label Operation for LDP Link Protection	557
Sample Multicast LDP Link Protection Configuration	563
Make-Before-Break	564
Caveats and Limitations	566
Understanding Multicast-Only Fast Reroute	566
PIM Functionality	569
Multipoint LDP Functionality	570
Packet Forwarding	571

Limitations and Caveats	572
Configuring Multicast-Only Fast Reroute	573
Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain	576
Example: Configuring LDP Downstream on Demand	592
Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs	597
Understanding Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs	598
How M-LDP Works	599
Terminology	603
Ingress Join Translation and Pseudo Interface Handling	604
Ingress Splicing	604
Reverse Path Forwarding	604
LSP Root Detection	604
Egress Join Translation and Pseudo Interface Handling	604
Egress Splicing	605
Supported Functionality	605
Unsupported Functionality	605
LDP Functionality	606
Egress LER Functionality	606
Transit LSR Functionality	606
Ingress LER Functionality	606
Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs	607
Configuring Miscellaneous LDP Properties	628
Configuring LDP to Use the IGP Route Metric	628
Preventing Addition of Ingress Routes to the inet.0 Routing Table	628
Multiple-Instance LDP and Carrier-of-Carriers VPNs	629
Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router	629
Enabling LDP over RSVP-Established LSPs	629
Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks ..	630
Configuring the TCP MD5 Signature for LDP Sessions	630
Configuring LDP Session Protection	631
Disabling SNMP Traps for LDP	632
Configuring LDP Synchronization with the IGP on LDP Links	632
Configuring LDP Synchronization with the IGP on the Router	633
Configuring the Label Withdrawal Timer	633
Ignoring the LDP Subnet Check	633
Configuring LDP LSP Traceroute	634
Collecting LDP Statistics	635
LDP Statistics Output	635
Disabling LDP Statistics on the Penultimate-Hop Router	636
LDP Statistics Limitations	637
Tracing LDP Protocol Traffic	637
Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels	637
Tracing LDP Protocol Traffic Within FECs	638

	Examples: Tracing LDP Protocol Traffic	639
Part 5	Configuring Circuit Cross-Connect (CCC) and Translational Cross-Connect (TCC)	
Chapter 15	CCC and TCC Overview	643
	CCC Overview	643
	Transmitting Nonstandard BPDUs	644
	TCC Overview	644
Chapter 16	Configuring CCC and TCC	647
	Configuring Layer 2 Switching Cross-Connects Using CCC	647
	Configuring the CCC Encapsulation for Layer 2 Switching	
	Cross-Connects	648
	Configuring ATM Encapsulation for Layer 2 Switching	
	Cross-Connects	648
	Configuring Ethernet Encapsulation for Layer 2 Switching	
	Cross-Connects	649
	Configuring Ethernet VLAN Encapsulation for Layer 2 Switching	
	Cross-Connects	649
	Configuring Aggregated Ethernet Encapsulation for Layer 2 Switching	
	Cross-Connects	650
	Configuring Frame Relay Encapsulation for Layer 2 Switching	
	Cross-Connects	651
	Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching	
	Cross-Connects	652
	Configuring the CCC Connection for Layer 2 Switching Cross-Connects	652
	Configuring MPLS for Layer 2 Switching Cross-Connects	652
	Example: Configuring a Layer 2 Switching Cross-Connect	653
	Configuring MPLS LSP Tunnel Cross-Connects Using CCC	655
	Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects	656
	Configuring the CCC Connection for LSP Tunnel Cross-Connects	657
	Example: Configuring an LSP Tunnel Cross-Connect	658
	Configuring TCC	659
	Configuring the Encapsulation for Layer 2 Switching TCCs	659
	Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching	
	TCCs	660
	Configuring ATM Encapsulation for Layer 2 Switching TCCs	660
	Configuring Frame Relay Encapsulation for Layer 2 Switching TCCs	660
	Configuring Ethernet Encapsulation for Layer 2 Switching TCCs	661
	Configuring Ethernet Extended VLAN Encapsulation for Layer 2	
	Switching TCCs	661
	Configuring ARP for Ethernet and Ethernet Extended VLAN	
	Encapsulations	662
	Configuring the Connection for Layer 2 Switching TCCs	663
	Configuring MPLS for Layer 2 Switching TCCs	663
	CCC and TCC Graceful Restart	664
	Configuring CCC and TCC Graceful Restart	665

	Configuring CCC Switching for Point-to-Multipoint LSPs	665
	Configuring the Point-to-Multipoint LSP Switch on Ingress PE Routers . . .	666
	Configuring Local Receivers on a Point-to-Multipoint CCC LSP Switch on Ingress PE Routers	666
	Configuring the Point-to-Multipoint LSP Switch on Egress PE Routers . . .	667
Part 6	Configuring GMPLS	
Chapter 17	GMPLS Overview	671
	Introduction to GMPLS	671
	GMPLS Terms and Acronyms	672
	Supported GMPLS Standards	673
	GMPLS Operation	674
	GMPLS and OSPF	675
	GMPLS and CSPF	675
	GMPLS Features	676
Chapter 18	Configuring GMPLS	677
	LMP Configuration Overview	677
	Configuring LMP Traffic Engineering Links	678
	Configuring the Local IP Address for Traffic Engineering Links	679
	Configuring the Remote IP Address for Traffic Engineering Links	679
	Configuring the Remote ID for Traffic Engineering Links	680
	Configuring LMP Peers	680
	Configuring the ID for LMP Peers	681
	Configuring the Interface for Control Channels Between LMP Peers	681
	Configuring the LMP Control Channel Interface for the Peer	681
	Configuring the Remote IP Address for LMP Control Channels	682
	Configuring Hello Message Intervals for LMP Control Channels	683
	Controlling Message Exchange for LMP Control Channels	684
	Preventing the Local Peer from Initiating LMP Negotiation	684
	Associating Traffic Engineering Links with LMP Peers	684
	Disabling the Traffic Engineering Link for LMP Peers	685
	Configuring RSVP and OSPF for LMP Peer Interfaces	685
	Configuring RSVP Signaling for LMP Peer Interfaces	685
	Configuring OSPF Routing for LMP Peer Interfaces	686
	Configuring the Hello Interval for LMP Peer Interfaces	686
	Configuring MPLS Paths for GMPLS	686
	Tracing LMP Traffic	687
	Configuring MPLS LSPs for GMPLS	688
	Configuring the Encoding Type	688
	Configuring the GPID	689
	Configuring the Signal Bandwidth Type	689
	Configuring GMPLS Bidirectional LSPs	689
	Allowing Nonpacket GMPLS LSPs to Establish Paths Through Routers Running Junos OS	690
	Gracefully Tearing Down GMPLS LSPs	690
	Temporarily Deleting GMPLS LSPs	690
	Permanently Deleting GMPLS LSPs	691

	Configuring the Graceful Deletion Timeout Interval	691
	GMPLS RSVP-TE VLAN LSP Signaling Overview	692
	Understanding GMPLS RSVP-TE Signaling	692
	Need for GMPLS RSVP-TE VLAN LSP Signaling	692
	GMPLS RSVP-TE VLAN LSP Signaling Functionality	694
	LSP Hierarchy with GMPLS RSVP-TE VLAN LSP	695
	Path Specification for GMPLS RSVP-TE VLAN LSP	695
	GMPLS RSVP-TE VLAN LSP Configuration	695
	Associated Bidirectional Packet LSP	696
	Make-Before-Break for Associated Bidirectional Packet and GMPLS RSVP-TE VLAN LSP	697
	Supported and Unsupported Features	698
	Example: Configuring GMPLS RSVP-TE VLAN LSP Signaling	698
Chapter 19	Using a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs over a Single RSVP LSP	723
	Hierarchy of RSVP LSPs Overview	723
	Hierarchy of RSVP LSPs Terminology	723
	Hierarchy of RSVP LSPs Standard	724
	Hierarchy of RSVP LSPs	724
	Advertising the Forwarding Adjacency with OSPF	724
	Configuring a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs Over a Single RSVP LSP	724
	Configuring an RSVP LSP on Ingress Routers	725
	Configuring Forwarding Adjacencies	725
	Configuring the Local IP Address for Forwarding Adjacencies	725
	Configuring the Remote IP Address for Forwarding Adjacencies	726
	Configuring the LSP for Forwarding Adjacencies	726
	Configuring RSVP for Forwarding Adjacencies	726
	Advertising Forwarding Adjacencies Using OSPF	727
Part 7	Configuring Path Computation Element Protocol (PCEP)	
Chapter 20	PCEP Overview	731
	PCEP Overview	731
Chapter 21	Configuring PCEP	733
	Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE	733
	Support of Path Computation Element Protocol for RSVP-TE Overview	733
	Understanding MPLS RSVP-TE	733
	Current MPLS RSVP-TE Limitations	735
	Use of an External Path Computing Entity	736
	Components of External Path Computing	737
	Interaction Between a PCE and a PCC Using PCEP	739
	LSP Behavior with External Computing	740
	Configuration Statements Supported for External Computing	742
	PCE-Controlled LSP Protection	742
	Auto-Bandwidth and PCE-Controlled LSP	742

	Impact of Client-Side PCE Implementation on Network Performance	743
	Example: Configuring the Path Computation Element Protocol for MPLS RSVP-TE	743
Part 8	Troubleshooting Information	
Chapter 22	Troubleshooting MPLS	759
	Verify MPLS Interfaces	759
	Verify That Node-Link Protection Is Up	761
	Verify That Link Protection Is Up	768
	Verify One-to-One Backup	772
	Verify That the Primary Path Is Operational	779
	Verify That the Secondary Path Is Established	780
	Checklist for Checking the MPLS Layer	782
	Checking the MPLS Layer	783
	Verify the LSP	785
	Verify the LSP Route on the Transit Router	788
	Verify the LSP Route on the Ingress Router	789
	Verify MPLS Labels with the traceroute Command	790
	Verify MPLS Labels with the ping Command	791
	Verify the MPLS Configuration	792
	Take Appropriate Action	794
	Verify the LSP Again	795
	Checklist for Working with the Layered MPLS Troubleshooting Model	798
	Understanding the Layered MPLS Troubleshooting Model	798
	Verify That Load Balancing Is Working	805
	Verify the Operation of Uneven Bandwidth Load Balancing	808
Part 9	Configuration Statements	
Chapter 23	MPLS Configuration Statements	813
	[edit protocols mpls] Hierarchy Level	817
	[edit logical-systems] Hierarchy Level	823
	[edit protocols connections] Hierarchy Level	824
	[edit protocols link-management] Hierarchy Level	824
	adaptive	825
	adjust-interval	826
	adjust-threshold	826
	adjust-threshold-activate-bandwidth	827
	adjust-threshold-overflow-limit	827
	adjust-threshold-underflow-limit	828
	admin-down	828
	admin-group (for Interfaces)	829
	admin-group (for LSPs)	829
	admin-group-extended	830
	admin-groups	831
	admin-groups-extended	832
	admin-groups-extended-range	833

advertise-mode (MPLS)	834
advertisement-hold-time	835
allow-fragmentation	835
always-mark-connection-protection-tlv	836
associate-backup-pe-groups	836
associate-lsp	837
auto-bandwidth (MPLS Tunnel)	838
auto-bandwidth (MPLS Statistics)	839
auto-policing	840
backup-pe-group	841
bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)	842
bandwidth (Static LSP)	843
bandwidth-model	844
bandwidth-percent	845
bfd-liveness-detection (Protocols MPLS)	846
class-of-service (Protocols MPLS)	847
container-label-switched-path	848
corouted-bidirectional	849
corouted-bidirectional-passive	849
credibility	850
database	851
delay (querier)	852
delay (responder)	853
description (Protocols MPLS)	854
deselect-on-bandwidth-failure	855
diffserv-te	856
disable (Protocols MPLS)	857
dynamic-tunnels	858
egress-protection (MPLS)	859
encoding-type	860
entropy-label	860
ethernet-vlan (Protocols Link Management)	861
exclude (for Administrative Groups)	861
exclude (for Fast Reroute)	862
exclude-srlg	863
expand-loose-hop	864
explicit-null (Protocols MPLS)	865
export (MPLS Traffic engineering database)	866
failure-action (Protocols MPLS)	867
family mpls	868
fast-reroute (Protocols MPLS)	870
fate-sharing	871
from (Protocols MPLS)	872
gpip	873
gre (Routing Options)	874
hop-limit	875
import (MPLS Traffic Engineering Database)	876
include-all (for Administrative Groups)	877
include-all (for Fast Reroute)	877

include-any (for Administrative Groups)	878
include-any (for Fast Reroute)	878
ingress (LSP)	879
install (Protocols MPLS)	880
ingress-policy	881
interface (Protocols MPLS)	882
inter-domain	883
ipv6-tunneling	883
label-switched-path (Protocols MPLS)	884
label-switched-path-template (Container LSP)	887
ldp-tunneling	887
least-fill	887
link-protection (Dynamic LSPs)	888
link-protection (Static LSPs)	889
load-balance-label-capability	889
log-updown (Protocols MPLS)	890
loss (querier)	891
loss (responder)	892
loss-delay (querier)	893
lsp-attributes	894
maximum-bandwidth (Protocols MPLS)	894
maximum-labels	895
minimum-bandwidth-adjust-interval	896
minimum-bandwidth-adjust-threshold-change	896
minimum-bandwidth-adjust-threshold-value	897
metric (Protocols MPLS)	898
minimum-bandwidth	898
monitor-bandwidth	899
most-fill	899
mpls (Protocols)	899
mpls-tp-mode	900
mtu-signaling	900
next-hop (Protocols MPLS)	901
no-bfd-triggered-local-repair	902
no-cspf	903
no-decrement-ttl	904
no-install-to-address	905
no-load-balance-label-capability	905
no-mcast-replication	906
no-propagate-ttl	907
no-transit-statistics	907
no-trap	908
node-protection (Static LSP)	909
normalization	910
oam (Protocols MPLS)	911
optimize-adaptive-teardown	912
optimize-aggressive	913
optimize-hold-dead-delay	914
optimize-switchover-delay	915

optimize-timer (Protocols MPLS)	916
p2mp (Protocols MPLS)	917
p2mp-lsp-next-hop	918
path (Protocols MPLS)	919
path-mtu	920
per-prefix-label	921
performance-monitoring (Protocols MPLS)	922
policing (Protocols MPLS)	923
pop	924
preference (Protocols MPLS)	925
primary (Protocols MPLS)	926
priority (Protocols MPLS)	927
protection-revert-time	928
push	929
random	930
record	931
retry-limit	932
retry-timer	932
revert-timer	933
rpf-check-policy (Routing Options)	934
rsvp-error-hold-time	935
sampling (Protocols MPLS)	936
secondary (Protocols MPLS)	937
select	938
signal-bandwidth	938
smart-optimize-timer	939
soft-preemption (Protocols MPLS)	940
splitting-merging	941
srlg	942
srlg-cost	943
srlg-value	943
standby	944
static-label-switched-path	945
statistics (Protocols MPLS)	947
swap	948
switch-away-lsps	949
switching-type	950
sync-active-path-bandwidth	951
te-class-matrix	952
to	953
traceoptions (Protocols MPLS)	954
traffic-class (delay)	956
traffic-class (loss)	958
traffic-class (loss-delay)	960
traffic-engineering (Protocols MPLS)	962
traffic-engineering (Protocols BGP)	963
transit-lsp-association	964
ultimate-hop-popping	965

Chapter 24	RSVP Configuration Statements	967
	[edit protocols rsvp] Hierarchy Level	968
	[edit protocols rsvp] Hierarchy Level	971
	admin-group	973
	aggregate (Protocols RSVP)	974
	authentication-key (Protocols RSVP)	975
	bandwidth (Protocols RSVP)	976
	bypass (Signaled LSP)	977
	bypass (Static LSP)	978
	class-of-service (Protocols RSVP)	979
	destination-networks	980
	devices	981
	disable (Protocols RSVP)	982
	dynamic-bidirectional-transport	983
	fast-reroute (Protocols RSVP)	983
	graceful-deletion-timeout	984
	graceful-restart (Protocols RSVP)	985
	hello-acknowledgements	986
	hello-interval (Protocols RSVP)	986
	hop-limit	987
	interface (Protocols RSVP)	988
	keep-multiplier	989
	label-switched-path-template (Multicast)	990
	link-protection (RSVP)	992
	load-balance (Protocols RSVP)	993
	max-bypasses	994
	no-local-reversion	995
	node-hello	996
	no-adjacency-down-notification (Protocols IS-IS)	997
	no-cspf (Protocols RSVP)	998
	no-interface-hello	998
	no-neighbor-down-notification	999
	no-node-id-subobject	999
	no-p2mp-sublsp	1000
	node-link-protection (Protocols MPLS)	1000
	optimize-timer (Protocols RSVP)	1001
	path (Protocols RSVP)	1002
	peer-interface (Protocols RSVP)	1003
	preemption	1004
	priority (Protocols RSVP)	1005
	refresh-time	1006
	reliable	1006
	rsvp	1007
	rsvp-te (Routing Options)	1008
	setup-protection	1008
	soft-preemption (Protocols RSVP)	1009
	static-label-switched-path	1010
	subscription	1011
	traceoptions (Protocols RSVP)	1012

Chapter 25

transit	1014
tunnel-services (RSVP)	1015
ultimate-hop-popping	1016
update-threshold	1017
LDP Configuration Statements	1019
[edit protocols ldp] Hierarchy Level	1021
allow-subnet-mismatch	1023
authentication-algorithm	1024
authentication-key (Protocols LDP)	1026
authentication-key-chain (Protocols LDP)	1027
auto-targeted-session	1028
bfd-liveness-detection (Protocols LDP)	1029
deaggregate	1030
disable (Protocols LDP)	1031
dod-request-policy	1032
downstream-on-demand	1032
ecmp	1033
egress-policy	1033
explicit-null (Protocols LDP)	1034
export (Protocols LDP)	1034
failure-action (Protocols LDP)	1035
fec	1036
graceful-restart (Protocols LDP)	1037
hello-interval (Protocols LDP)	1038
helper-disable (LDP)	1039
holddown-interval	1040
hold-time (Protocols LDP)	1041
ignore-lsp-metrics	1042
igp-synchronization	1042
import (Protocols LDP)	1043
ingress-policy	1044
interface (Protocols LDP)	1045
keepalive-interval	1046
keepalive-timeout	1047
l2-smart-policy	1047
label-withdrawal-delay	1048
ldp	1049
ldp-synchronization	1052
log-updown (Protocols LDP)	1053
make-before-break (LDP)	1054
maximum-neighbor-recovery-time	1055
mldp-inband-signalling (Protocols Multipoint LDP)	1056
mofrr-asm-starg (Multicast-Only Fast Reroute in a PIM Domain)	1057
mofrr-disjoint-upstream-only (Multicast-Only Fast Reroute in a PIM Domain)	1058
mofrr-no-backup-join (Multicast-Only Fast Reroute in a PIM Domain)	1059
mofrr-primary-selection-by-routing (Multicast-Only Fast Reroute)	1060
no-forwarding	1061

	oam (Protocols LDP)	1062
	p2mp (Protocols LDP)	1063
	p2mp-ldp-next-hop	1064
	periodic-traceroute	1065
	policing (Protocols LDP)	1067
	policy (Multicast-Only Fast Reroute)	1068
	policy (Protocols Multipoint LDP)	1070
	preference (Protocols LDP)	1071
	reconnect-time	1072
	recovery-time	1073
	session (ldp)	1074
	session-protection	1075
	stream-protection (Multicast-Only Fast Reroute)	1076
	strict-targeted-hellos	1077
	targeted-hello	1077
	traceoptions (Protocols LDP)	1078
	track-igp-metric	1080
	traffic-statistics (Protocols LDP)	1081
	transport-address	1083
	version (BFD)	1084
Chapter 26	CCC and TCC Configuration Statements	1085
	connections (Circuits)	1086
	encapsulation (Logical Interface)	1087
	encapsulation (Physical Interface)	1091
	interface-switch	1096
	lsp-switch	1097
	output-interface (CCC)	1097
	p2mp-receive-switch	1098
	p2mp-transmit-switch	1099
	remote-interface-switch	1100
Chapter 27	GMPLS Configuration Statements	1101
	address (Peer)	1102
	control-channel (Protocols Link Management Peer)	1102
	dead-interval	1103
	disable (GMPLS)	1104
	disable (OSPF)	1105
	hello-dead-interval	1106
	hello-interval (LMP)	1107
	hello-interval (Protocols OSPF)	1108
	interface (Protocols Link Management)	1109
	label-switched-path (Protocols Link Management)	1109
	link-management	1110
	lmp-control-channel	1111
	lmp-protocol	1111
	local-address (Protocols Link Management)	1112
	passive (Protocols Link Management)	1112
	peer (Protocols LMP)	1113
	peer-interface (Protocols OSPF)	1114

	remote-address (for LMP Control Channel)	1115
	remote-address (for LMP Traffic Engineering)	1115
	remote-id	1116
	retransmission-interval	1116
	retransmit-interval (OSPF)	1117
	retry-limit (Protocols Link Management)	1118
	te-link	1119
	traceoptions (Protocols Link Management)	1120
	transit-delay (OSPF)	1122
	upstream-label	1123
Chapter 28	PCEP Configuration Statements	1125
	[edit protocols pcep] Hierarchy Level	1125
	pcep	1127
	delegation-cleanup-timeout	1128
	delegation-priority	1128
	destination-ipv4-address	1129
	destination-port	1129
	lsp-external-controller	1130
	max-unknown-messages	1130
	message-rate-limit	1131
	pce	1132
	pce-group (PCE)	1133
	pce-group (Protocols PCEP)	1134
	pce-type	1135
	querier (performance-monitoring)	1136
	traceoptions (PCE)	1137
	traceoptions (Protocols PCEP)	1139
	update-rate-limit	1140
Part 10	Operational Commands	
Chapter 29	MPLS Operational Commands	1143
	clear mpls lsp	1145
	clear mpls container-lsp	1147
	clear performance-monitoring mpls lsp	1149
	monitor mpls delay rsvp	1150
	monitor mpls loss rsvp	1154
	monitor mpls loss-delay rsvp	1159
	ping mpls bgp	1163
	ping mpls lsp-end-point	1165
	request mpls container-lsp	1167
	request mpls lsp adjust-autobandwidth	1168
	show connections	1170
	show link-management	1173
	show link-management peer	1177
	show link-management routing	1179
	show link-management statistics	1182
	show link-management te-link	1184
	show mpls admin-groups	1186

	show mpls call-admission-control	1188
	show mpls container-lsp	1190
	show mpls context-identifier	1197
	show mpls cspf	1199
	show mpls diffserv-te	1201
	show mpls egress-protection	1203
	show mpls interface	1205
	show mpls label usage	1207
	show mpls lsp	1209
	show mpls lsp autobandwidth	1227
	show mpls path	1230
	show mpls srlg	1232
	show mpls static-lsp	1233
	show performance-monitoring mpls lsp	1236
	show ted database	1242
	show ted link	1250
	show ted protocol	1253
	traceroute mpls bgp	1255
Chapter 30	RSVP Operational Commands	1259
	clear rsvp session	1260
	clear rsvp statistics	1262
	monitor label-switched-path	1263
	ping mpls rsvp	1266
	show rsvp interface	1271
	show rsvp neighbor	1276
	show rsvp session	1281
	show rsvp statistics	1291
	show rsvp version	1295
	traceroute mpls rsvp	1298
Chapter 31	LDP Operational Commands	1303
	clear ldp neighbor	1304
	clear ldp session	1305
	clear ldp statistics	1306
	ping mpls ldp	1307
	show ldp database	1310
	show ldp fec-filters	1319
	show ldp interface	1320
	show ldp neighbor	1322
	show ldp overview	1324
	show ldp p2mp tunnel	1328
	show ldp path	1329
	show ldp route	1331
	show ldp session	1335
	show ldp statistics	1341
	show ldp traffic-statistics	1345
	show security keychain	1349
	traceroute mpls ldp	1352

Chapter 32	CCC and TCC Operational Commands	1355
	show connections	1356
	show route ccc	1359
	show route forwarding-table	1360
Chapter 33	PCEP Operational Commands	1375
	clear path-computation-client statistics	1376
	request path-computation-client active-pce	1377
	show path-computation-client active-pce	1378
	show path-computation-client statistics	1382
Part 11	Index	
	Index	1389

List of Figures

Part 1	Understanding Traffic Engineering	
Chapter 1	Traffic Engineering Overview	3
	Figure 1: Junos OS Implementation of BGP Link-State Distribution	11
Part 2	Configuring MPLS	
Chapter 2	MPLS Overview	19
	Figure 2: Label Encoding	26
	Figure 3: Class-of-Service Bits	27
	Figure 4: CSPF Computation Process	31
	Figure 5: Aggregation Router A Dual-Homed on Core Routers B and C	34
	Figure 6: Typical SPF Tree, Sourced from Router A	35
	Figure 7: Modified SPF Tree, Using LSP A–D as a Shortcut	35
	Figure 8: Modified SPF Tree, Using LSP A–D and LSP A–E as Shortcuts	36
	Figure 9: IGP Shortcuts	37
	Figure 10: IGP Shortcuts in a Bigger Network	37
	Figure 11: SPF Computations with Advertised LSPs	38
	Figure 12: MPLS Application Topology	41
	Figure 13: How BGP Determines How to Reach Next-Hop Addresses	41
	Figure 14: Routing and Forwarding Tables, traffic-engineering bgp	44
	Figure 15: Routing and Forwarding Tables, traffic-engineering bgp-igp	45
	Figure 16: Detours Established for an LSP Using Fast Reroute	47
	Figure 17: Detour After the Link from Router B to Router C Fails	47
	Figure 18: Detours Merging into Other Detours	48
	Figure 19: Basic Bidirectional Measurement	55
Chapter 3	Configuring MPLS Routers	59
	Figure 20: IPv6 Networks Linked by MPLS IPv4 Tunnels	72
	Figure 21: MPLS-TP OAM Associated Bidirectional LSPs	139
	Figure 22: MPLS Inter-AS Link-Node Protection Conceptual Topology	149
	Figure 23: MPLS Inter-AS Link-Node Protection Example Topology	151
	Figure 24: Egress Protection LSP Configured from Router PE1 to Router PE2	164
	Figure 25: Egress Protection LSP Configured from Router PE1 to Router PE2	169
	Figure 26: Co-located PLR and protector in collocated protector model	185
Chapter 4	Configuring MPLS-Signaled LSPs	211
	Figure 27: Corouted Bidirectional LSP	220
	Figure 28: Penultimate-Hop Popping for an LSP	222
	Figure 29: Ultimate-Hop Popping for an LSP	222
	Figure 30: least-fill Load Balancing Algorithm Example	253
Chapter 5	Configuring Static and Explicit-Path LSPs	271

	Figure 31: Static MPLS Configuration	273
Chapter 6	Configuring Point-to-Multipoint LSPs	281
	Figure 32: Point-to-Multipoint LSPs	282
	Figure 33: RSVP-Signaled Point-to-Multipoint LSP	286
Chapter 8	Configuring Miscellaneous MPLS Properties	333
	Figure 34: Topology with BFD-Triggered Local Repair	358
	Figure 35: Link-State Distribution Using BGP	365
	Figure 36: Sample Topology	384
	Figure 37: Dynamic Bandwidth Management Using Container LSP	414
	Figure 38: Configuring On-Demand Loss and Delay Measurement	440
	Figure 39: Configuring Pro-Active Loss and Delay Measurements	451
Part 3	Configuring RSVP	
Chapter 10	Configuring RSVP	471
	Figure 40: Penultimate-Hop Popping for an LSP	486
	Figure 41: Ultimate-Hop Popping for an LSP	486
Chapter 11	Configuring RSVP Link Protection and Node Protection to Protect from Traffic Failures	495
	Figure 42: Link Protection Creating a Bypass LSP for the Protected Interface . .	496
	Figure 43: Node Protection Creating a Next-Next-Hop Bypass LSP	497
	Figure 44: One-to-One Backup	500
	Figure 45: Facility Backup	501
Part 4	Configuring LDP	
Chapter 13	LDP Overview	519
	Figure 46: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs	524
	Figure 47: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs . .	524
Chapter 14	Configuring LDP	527
	Figure 48: Multicast LDP Support for Targeted LDP Session	551
	Figure 49: Incomplete Coverage Problem with LFA	552
	Figure 50: Manually Configured RSVP LSP Coverage	553
	Figure 51: Dynamically Configured RSVP LSP Coverage	554
	Figure 52: Multicast LDP Link Protection Sample Topology	556
	Figure 53: LDP Label Operation Sample Topology	557
	Figure 54: Unicast LDP Label Operation	559
	Figure 55: Multicast LDP Label Operation	560
	Figure 56: LDP Link Protection Label Operation	563
	Figure 57: MoFRR Sample Topology	568
	Figure 58: MoFRR IP Route Lookup in the Packet Forwarding Engine	571
	Figure 59: MoFRR MPLS Route Lookup in the Packet Forwarding Engine	572
	Figure 60: MoFRR in a Multipoint LDP Domain	577
	Figure 61: Label Bindings in M-LDP Signaling	599
	Figure 62: Sample M-LDP Topology in PIM-Free MPLS Core	600
	Figure 63: Sample M-LDP Topology in PIM-Enabled MPLS Core	601

	Figure 64: M-LDP In-Band Signaling for Point-to-Multipoint LSPs Example Topology	608
Part 5	Configuring Circuit Cross-Connect (CCC) and Translational Cross-Connect (TCC)	
Chapter 15	CCC and TCC Overview	643
	Figure 65: TCC Example	644
Chapter 16	Configuring CCC and TCC	647
	Figure 66: Layer 2 Switching Cross-Connect	647
	Figure 67: Topology of a Frame Relay Layer 2 Switching Cross-Connect	653
	Figure 68: Sample Topology of a VLAN Layer 2 Switching Cross-Connect	654
	Figure 69: MPLS Tunnel Cross-Connect	655
	Figure 70: Example Topology of MPLS LSP Tunnel Cross-Connect	658
	Figure 71: Remote Interface Switch Connecting Two CE Routers Using CCC	664
Part 6	Configuring GMPLS	
Chapter 18	Configuring GMPLS	677
	Figure 72: Traditional Layer 2 Point-to-Point Services	693
	Figure 73: GMPLS RSVP-TE VLAN LSP	694
	Figure 74: Setting Up a GMPLS VLAN LSP	700
	Figure 75: Data Traffic Flow of GMPLS VLAN LSP	704
	Figure 76: Configuring GMPLS RSVP-TE VLAN LSP Signaling	705
Part 7	Configuring Path Computation Element Protocol (PCEP)	
Chapter 20	PCEP Overview	731
	Figure 77: PCEP Session	731
Chapter 21	Configuring PCEP	733
	Figure 78: Example MPLS Traffic Engineering	736
	Figure 79: PCC and RSVP-TE	739
	Figure 80: Example PCE for MPLS RSVP-TE	740
	Figure 81: Configuring PCEP for MPLS RSVP-TE	745
Part 8	Troubleshooting Information	
Chapter 22	Troubleshooting MPLS	759
	Figure 82: Checking the MPLS Layer	784
	Figure 83: MPLS Network Broken at the MPLS Layer	784
	Figure 84: Layered MPLS Network Troubleshooting Model	799
	Figure 85: MPLS Basic Network Topology Example	801

List of Tables

	About the Documentation	xxxvii
	Table 1: Notice Icons	xxxix
	Table 2: Text and Syntax Conventions	xl
Part 2	Configuring MPLS	
Chapter 4	Configuring MPLS-Signaled LSPs	211
	Table 3: MPLS LSP Load Balancing Options	234
	Table 4: MPLS CoS Values	246
Chapter 7	Configuring DiffServ-Aware Traffic Engineering to Achieve Service Level Guarantees on an MPLS network	309
	Table 5: Default Values for the Traffic Engineering Class Matrix	317
Chapter 8	Configuring Miscellaneous MPLS Properties	333
	Table 6: Sample Scenarios for Using 3, 4, or 5 MPLS Labels	334
	Table 7: LSP Sequence Order for Bin Packing	384
	Table 8: LSP Sequence Order for Deadlock	385
	Table 9: LSP Sequence Order for Predictability	386
	Table 10: LSP Sequence Order for Predictability	386
	Table 11: Normalization with Per-LSP Autobandwidth Adjustment Changes	396
	Table 12: Normalization with Traffic Growth	397
	Table 13: Applicability of RSVP LSPs Configuration to a Container LSP	404
Part 3	Configuring RSVP	
Chapter 11	Configuring RSVP Link Protection and Node Protection to Protect from Traffic Failures	495
	Table 14: One-to-One Backup Compared with Facility Backup	499
Part 4	Configuring LDP	
Chapter 14	Configuring LDP	527
	Table 15: from Operators That Apply to LDP Received-Label Filtering	536
	Table 16: to Operators for LDP Outbound-Label Filtering	538
Part 7	Configuring Path Computation Element Protocol (PCEP)	
Chapter 21	Configuring PCEP	733
	Table 17: Applicability of MPLS and Existing LSP Configurations to a PCE-Controlled LSP	742

Part 8	Troubleshooting Information
Chapter 22	Troubleshooting MPLS 759
	Table 18: Checklist for Checking the MPLS Layer 783
	Table 19: Checklist for Working with the Layered MPLS Troubleshooting Model 798
Part 10	Operational Commands
Chapter 29	MPLS Operational Commands 1143
	Table 20: monitor mpls delay rsvp Output Fields 1151
	Table 21: monitor mpls loss rsvp Output Fields 1155
	Table 22: show connections Output Fields 1171
	Table 23: show link-management Output Fields 1173
	Table 24: show link-management peer Output Fields 1177
	Table 25: show link-management routing Output Fields 1179
	Table 26: show link-management statistics Output Fields 1182
	Table 27: show link-management te-link Output Fields 1184
	Table 28: show mpls admin-groups Output Fields 1186
	Table 29: show mpls call-admission-control Output Fields 1188
	Table 30: show mpls container-lsp Output Fields 1191
	Table 31: show mpls lsp Output Fields 1197
	Table 32: show mpls cspf Output Fields 1199
	Table 33: show mpls diffserv-te Output Fields 1201
	Table 34: show mpls lsp Output Fields 1203
	Table 35: show mpls interface Output Fields 1205
	Table 36: show mpls label usage Fields 1207
	Table 37: show mpls lsp Output Fields 1211
	Table 38: show mpls lsp autobandwidth Output Fields 1227
	Table 39: show mpls path Output Fields 1230
	Table 40: show mpls srlg Output Fields 1232
	Table 41: show mpls static-lsp Output Fields 1234
	Table 42: show performance-monitoring mpls lsp Output Fields 1237
	Table 43: show ted database Output Fields 1242
	Table 44: show ted link Output Fields 1250
	Table 45: show ted protocol Output Fields 1253
	Table 46: traceroute mpls bgp Output Fields 1256
Chapter 30	RSVP Operational Commands 1259
	Table 47: Output Control Keys for the monitor label-switched-path Command 1263
	Table 48: monitor label-switched-path Output Fields 1264
	Table 49: show rsvp interface Output Fields 1272
	Table 50: show rsvp neighbor Output Fields 1276
	Table 51: show rsvp session Output Fields 1283
	Table 52: show rsvp statistics Output Fields 1291
	Table 53: show rsvp version Output Fields 1295
	Table 54: traceroute mpls rsvp Output Fields 1299
Chapter 31	LDP Operational Commands 1303

	Table 55: show ldp database Output Fields	1311
	Table 56: show ldp fec-filters Output Fields	1319
	Table 57: show ldp interface Output Fields	1320
	Table 58: show ldp neighbor Output Fields	1322
	Table 59: show ldp overview Output Fields	1324
	Table 60: show ldp path Output Fields	1329
	Table 61: show ldp route Output Fields	1331
	Table 62: show ldp session Output Fields	1335
	Table 63: show ldp statistics Output Fields	1341
	Table 64: show ldp traffic-statistics Output Fields	1346
	Table 65: show security keychain Output Fields	1349
	Table 66: traceroute mpls ldp Output Fields	1353
Chapter 32	CCC and TCC Operational Commands	1355
	Table 67: show connections Output Fields	1357
	Table 68: show route forwarding-table Output Fields	1363
Chapter 33	PCEP Operational Commands	1375
	Table 69: show path-computation-client active-pce Output Fields	1378
	Table 70: show path-computation-client statistics Output Fields	1382

About the Documentation

- [Documentation and Release Notes on page xxxvii](#)
- [Supported Platforms on page xxxvii](#)
- [Using the Examples in This Manual on page xxxvii](#)
- [Documentation Conventions on page xxxix](#)
- [Documentation Feedback on page xli](#)
- [Requesting Technical Support on page xli](#)

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

Supported Platforms

For the features described in this document, the following platforms are supported:

- [ACX Series](#)
- [T Series](#)
- [MX Series](#)
- [M Series](#)
- [PTX Series](#)

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming

configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xsl;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xml; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see the *CLI User Guide*.

Documentation Conventions

Table 1 on page xxxix defines notice icons used in this guide.

Table 1: Notice Icons







Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xl defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric metric>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (string1 string2 string3)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [community-ids]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page at the Juniper Networks Technical Documentation site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.

PART 1

Understanding Traffic Engineering

- [Traffic Engineering Overview on page 3](#)

CHAPTER 1

Traffic Engineering Overview

- [Traffic Engineering Capabilities on page 3](#)
- [Components of Traffic Engineering on page 4](#)
- [Packet Forwarding Component on page 4](#)
- [Packet Forwarding Based on Label Swapping on page 4](#)
- [How a Packet Traverses an MPLS Backbone on page 5](#)
- [Information Distribution Component on page 5](#)
- [Path Selection Component on page 6](#)
- [Offline Path Planning and Analysis on page 6](#)
- [Signaling Component on page 7](#)
- [Flexible LSP Calculation and Configuration on page 7](#)
- [Link-State Distribution Using BGP Overview on page 8](#)

Traffic Engineering Capabilities

The task of mapping traffic flows onto an existing physical topology is called *traffic engineering*. Traffic engineering provides the ability to move traffic flow away from the shortest path selected by the interior gateway protocol (IGP) and onto a potentially less congested physical path across a network.

Traffic engineering provides the capabilities to do the following:

- Route primary paths around known bottlenecks or points of congestion in the network.
- Provide precise control over how traffic is rerouted when the primary path is faced with single or multiple failures.
- Provide more efficient use of available aggregate bandwidth and long-haul fiber by ensuring that subsets of the network do not become overutilized while other subsets of the network along potential alternate paths are underutilized.
- Maximize operational efficiency.
- Enhance the traffic-oriented performance characteristics of the network by minimizing packet loss, minimizing prolonged periods of congestion, and maximizing throughput.
- Enhance statistically bound performance characteristics of the network (such as loss ratio, delay variation, and transfer delay) required to support a multiservices Internet.

Components of Traffic Engineering

In the Junos[®] operating system (OS), traffic engineering is implemented with MPLS and RSVP. Traffic engineering is composed of four functional components:

- [Packet Forwarding Component on page 4](#)
- [Information Distribution Component on page 5](#)
- [Path Selection Component on page 6](#)
- [Signaling Component on page 7](#)

Packet Forwarding Component

The packet forwarding component of the Junos traffic engineering architecture is MPLS, which is responsible for directing a flow of IP packets along a predetermined path across a network. This path is called a *label-switched path (LSP)*. LSPs are simplex; that is, the traffic flows in one direction from the head-end (ingress) router to a tail-end (egress) router. Duplex traffic requires two LSPs: one LSP to carry traffic in each direction. An LSP is created by the concatenation of one or more label-switched hops, allowing a packet to be forwarded from one router to another across the MPLS domain.

When an ingress router receives an IP packet, it adds an MPLS header to the packet and forwards it to the next router in the LSP. The labeled packet is forwarded along the LSP by each router until it reaches the tail end of the LSP, the egress router. At this point the MPLS header is removed, and the packet is forwarded based on Layer 3 information such as the IP destination address. The value of this scheme is that the physical path of the LSP is not limited to what the IGP would choose as the shortest path to reach the destination IP address.

This section discusses the following topics:

- [Packet Forwarding Based on Label Swapping on page 4](#)
- [How a Packet Traverses an MPLS Backbone on page 5](#)

Related Documentation

- [Components of Traffic Engineering on page 4](#)

Packet Forwarding Based on Label Swapping

The packet forwarding process at each router is based on the concept of label swapping. This concept is similar to what occurs at each Asynchronous Transfer Mode (ATM) switch in a permanent virtual circuit (PVC). Each MPLS packet carries a 4-byte encapsulation header that contains a 20-bit, fixed-length label field. When a packet containing a label arrives at a router, the router examines the label and copies it as an index to its MPLS forwarding table. Each entry in the forwarding table contains an interface-inbound label pair mapped to a set of forwarding information that is applied to all packets arriving on the specific interface with the same inbound label.

How a Packet Traverses an MPLS Backbone

This section describes how an IP packet is processed as it traverses an MPLS backbone network.

At the entry edge of the MPLS backbone, the IP header is examined by the ingress router. Based on this analysis, the packet is classified, assigned a label, encapsulated in an MPLS header, and forwarded toward the next hop in the LSP. MPLS provides a high degree of flexibility in the way that an IP packet can be assigned to an LSP. For example, in the Junos traffic engineering implementation, all packets arriving at the ingress router that are destined to exit the MPLS domain at the same egress router are forwarded along the same LSP.

Once the packet begins to traverse the LSP, each router uses the label to make the forwarding decision. The MPLS forwarding decision is made independently of the original IP header: the incoming interface and label are used as lookup keys into the MPLS forwarding table. The old label is replaced with a new label, and the packet is forwarded to the next hop along the LSP. This process is repeated at each router in the LSP until the packet reaches the egress router.

When the packet arrives at the egress router, the label is removed and the packet exits the MPLS domain. The packet is then forwarded based on the destination IP address contained in the packet's original IP header according to the traditional shortest path calculated by the IP routing protocol.

Information Distribution Component

Traffic engineering requires detailed knowledge about the network topology as well as dynamic information about network loading. To implement the information distribution component, simple extensions to the IGP are defined. Link attributes are included as part of each router's link-state advertisement. IS-IS extensions include the definition of new type length values (TLVs), whereas OSPF extensions are implemented with opaque link-state advertisements (LSAs). The standard flooding algorithm used by the link-state IGPs ensures that link attributes are distributed to all routers in the routing domain. Some of the traffic engineering extensions to be added to the IGP link-state advertisement include maximum link bandwidth, maximum reserved link bandwidth, current bandwidth reservation, and link coloring.

Each router maintains network link attributes and topology information in a specialized traffic engineering database. The traffic engineering database is used exclusively for calculating explicit paths for the placement of LSPs across the physical topology. A separate database is maintained so that the subsequent traffic engineering computation is independent of the IGP and the IGP's link-state database. Meanwhile, the IGP continues its operation without modification, performing the traditional shortest-path calculation based on information contained in the router's link-state database.

Related Documentation

- [Components of Traffic Engineering on page 4](#)

Path Selection Component

After network link attributes and topology information are flooded by the IGP and placed in the traffic engineering database, each ingress router uses the traffic engineering database to calculate the paths for its own set of LSPs across the routing domain. The path for each LSP can be represented by either a strict or loose explicit route. An explicit route is a preconfigured sequence of routers that should be part of the physical path of the LSP. If the ingress router specifies all the routers in the LSP, the LSP is said to be identified by a strict explicit route. If the ingress router specifies only some of the routers in the LSP, the LSP is described as a loose explicit route. Support for strict and loose explicit routes allows the path selection process to be given broad latitude whenever possible, but to be constrained when necessary.

The ingress router determines the physical path for each LSP by applying a Constrained Shortest Path First (CSPF) algorithm to the information in the traffic engineering database. CSPF is a shortest-path-first algorithm that has been modified to take into account specific restrictions when the shortest path across the network is calculated. Input into the CSPF algorithm includes:

- Topology link-state information learned from the IGP and maintained in the traffic engineering database
- Attributes associated with the state of network resources (such as total link bandwidth, reserved link bandwidth, available link bandwidth, and link color) that are carried by IGP extensions and stored in the traffic engineering database
- Administrative attributes required to support traffic traversing the proposed LSP (such as bandwidth requirements, maximum hop count, and administrative policy requirements) that are obtained from user configuration

As CSPF considers each candidate node and link for a new LSP, it either accepts or rejects a specific path component based on resource availability or whether selecting the component violates user policy constraints. The output of the CSPF calculation is an explicit route consisting of a sequence of router addresses that provides the shortest path through the network that meets the constraints. This explicit route is then passed to the signaling component, which establishes the forwarding state in the routers along the LSP.

Related Documentation

- [Components of Traffic Engineering on page 4](#)

Offline Path Planning and Analysis

Despite the reduced management effort resulting from online path calculation, an offline planning and analysis tool is still required to optimize traffic engineering globally. Online calculation takes resource constraints into account and calculates one LSP at a time. The challenge with this approach is that it is not deterministic. The order in which LSPs are calculated plays a critical role in determining each LSP's physical path across the network. LSPs that are calculated early in the process have more resources available to them than LSPs calculated later in the process because previously calculated LSPs

consume network resources. If the order in which the LSPs are calculated is changed, the resulting set of physical paths for the LSPs also can change.

An offline planning and analysis tool simultaneously examines each link's resource constraints and the requirements of each LSP. Although the offline approach can take several hours to complete, it performs global calculations, compares the results of each calculation, and then selects the best solution for the network as a whole. The output of the offline calculation is a set of LSPs that optimizes utilization of network resources. After the offline calculation is completed, the LSPs can be established in any order because each is installed according to the rules for the globally optimized solution.

Signaling Component

An LSP is not known to be workable until it is actually established by the signaling component. The signaling component, which is responsible for establishing LSP state and distributing labels, relies on a number of extensions to RSVP:

- The Explicit Route object allows an RSVP path message to traverse an explicit sequence of routers that is independent of conventional shortest-path IP routing. The explicit route can be either strict or loose.
- The Label Request object permits the RSVP path message to request that intermediate routers provide a label binding for the LSP that it is establishing.
- The Label object allows RSVP to support the distribution of labels without changing its existing mechanisms. Because the RSVP Resv message follows the reverse path of the RSVP path message, the Label object supports the distribution of labels from downstream nodes to upstream nodes.

Related Documentation

- [Components of Traffic Engineering on page 4](#)

Flexible LSP Calculation and Configuration

Traffic engineering involves mapping traffic flow onto a physical topology. You can determine the paths online using constraint-based routing. Regardless of how the physical path is calculated, the forwarding state is installed across the network through RSVP.

The Junos OS supports the following ways to route and configure an LSP:

- You can calculate the full path for the LSP offline and individually configure each router in the LSP with the necessary static forwarding state. This is analogous to the way some Internet service providers (ISPs) configure their IP-over-ATM cores.
- You can calculate the full path for the LSP offline and statically configure the ingress router with the full path. The ingress router then uses RSVP as a dynamic signaling protocol to install a forwarding state in each router along the LSP.
- You can rely on constraint-based routing to perform dynamic online LSP calculation. You configure the constraints for each LSP; then the network itself determines the path that best meets those constraints. Specifically, the ingress router calculates the entire LSP based on the constraints and then initiates signaling across the network.

- You can calculate a partial path for an LSP offline and statically configure the ingress router with a subset of the routers in the path; then you can permit online calculation to determine the complete path.

For example, consider a topology that includes two east-west paths across the United States: one in the north through Chicago and one in the south through Dallas. If you want to establish an LSP between a router in New York and one in San Francisco, you can configure the partial path for the LSP to include a single loose-routed hop of a router in Dallas. The result is an LSP routed along the southern path. The ingress router uses CSPF to compute the complete path and RSVP to install the forwarding state along the LSP.

- You can configure the ingress router with no constraints whatsoever. In this case, normal IGP shortest-path routing is used to determine the path of the LSP. This configuration does not provide any value in terms of traffic engineering. However, it is easy and might be useful in situations when services such as virtual private networks (VPNs) are needed.

In all these cases, you can specify any number of LSPs as backups for the primary LSP, thus allowing you to combine more than one configuration approach. For example, you might explicitly compute the primary path offline, set the secondary path to be constraint-based, and have the tertiary path be unconstrained. If a circuit on which the primary LSP is routed fails, the ingress router notices the outage from error notifications received from a downstream router or by the expiration of RSVP soft-state information. Then the router dynamically forwards traffic to a hot-standby LSP or calls on RSVP to create a forwarding state for a new backup LSP.

Link-State Distribution Using BGP Overview

- [Role of an Interior Gateway Protocol on page 8](#)
- [Limitations of an Interior Gateway Protocol on page 9](#)
- [Need for Spanning Link-State Distribution on page 9](#)
- [Using BGP as a Solution on page 10](#)
- [Supported and Unsupported Features on page 15](#)

Role of an Interior Gateway Protocol

An interior gateway protocol (IGP) is a type of protocol used for exchanging routing information between devices within an autonomous system (AS). Based on the method of computing the best path to a destination, the IGPs are divided into two categories:

- Link-state protocols—Advertise information about the network topology (directly connected links and the state of those links) to all routers using multicast addresses and triggered routing updates until all the routers running the link-state protocol have identical information about the internetwork. The best path to a destination is calculated based on constraints such as maximum delay, minimum available bandwidth, and resource class affinity.

OSPF and IS-IS are examples of link-state protocols.

- Distance vector protocols—Advertise complete routing table information to directly connected neighbors using a broadcast address. The best path is calculated based on the number of hops to the destination network.

RIP is an example of a distance vector protocol.

As the name implies, the role of an IGP is to provide routing connectivity within or internal to a given routing domain. A routing domain is a set of routers under common administrative control that share a common routing protocol. An AS can consist of multiple routing domains, where IGP functions to advertise and learn network prefixes (routes) from neighboring routers to build a route table that ultimately contains entries for all sources advertising reachability for a given prefix. IGP executes a route selection algorithm to select the best path between the local router and each destination, and provides full connectivity among the routers making up a routing domain.

In addition to advertising internal network reachability, IGPs are often used to advertise routing information that is external to that IGP's routing domain through a process known as route redistribution. Route redistribution is the process of exchanging routing information among distinct routing protocols to tie multiple routing domains together when intra-AS connectivity is desired.

Limitations of an Interior Gateway Protocol

While each individual IGP has its own advantages and limitations, the biggest limitations of IGP in general are performance and scalability.

IGPs are designed to handle the task of acquiring and distributing network topology information for traffic engineering purposes. While this model has served well, IGPs have inherent scaling limitations when it comes to distributing large databases. IGPs can autodetect neighbors, with which they acquire intra-area network topology information. However, the link-state database or a traffic engineering database has the scope of a single area or AS, thereby limiting applications, such as end-to-end traffic engineering, the benefit of having external visibility to make better decisions.

For label-switched networks, such as MPLS and Generalized MPLS (GMPLS), most existing traffic engineering solutions work in a single routing domain. These solutions do not work when a route from the ingress node to the egress node leaves the routing area or AS of the ingress node. In such cases, the path computation problem becomes complicated because of the unavailability of the complete routing information throughout the network. This is because service providers usually choose not to leak routing information beyond the routing area or AS for scalability constraints and confidentiality concerns.

Need for Spanning Link-State Distribution

One of the limitations of IGP is its inability to span link-state distribution outside a single area or AS. However, spanning link-state information acquired by an IGP across multiple areas or ASs has the following needs:

- LSP path computation—This information is used to compute the path for MPLS LSPs across multiple routing domains, for example an inter-area TE LSP.

- External path computing entities—External path computing entities, such as Application Layer Traffic Optimization (ALTO) and Path Computation Elements (PCE), perform path computations based on the network topology and current state of connections within the network, including traffic engineering information. This information is typically distributed by IGPs within the network.

However, because the external path computing entities cannot extract this information from the IGPs, they perform network monitoring to optimize network services.

Using BGP as a Solution

- [Overview on page 10](#)
- [Implementation on page 11](#)

Overview

To meet the needs for spanning link-state distribution across multiple domains, an exterior gateway protocol (EGP) is required to collect link-state and traffic engineering information from an IGP area, share it with external component, and use it for computing paths for interdomain MPLS LSPs.

BGP is a standardized EGP designed to exchange routing and reachability information between autonomous systems (ASs). BGP is a proven protocol that has better scaling properties because it can distribute millions of entries (for example, VPN prefixes) in a scalable fashion. BGP is the only routing protocol in use today that is suited to carry all of the routes in the Internet. This is largely because BGP runs on top of TCP and can make use of TCP flow control. In contrast, the internal gateway protocols (IGPs) do not have flow control. When IGPs have too much route information, they begin to churn. When BGP has a neighboring speaker that is sending information too quickly, BGP can throttle down the neighbor by delaying TCP acknowledgments.

Another benefit of BGP is that it uses type, length, value (TLV) tuples and network layer reachability information (NLRI) that provide seemingly endless extensibility without the need for the underlying protocol to be altered.

The distribution of link-state information across domains is regulated using policies to protect the interests of the service provider. This requires a control over the topology distribution using policies. BGP with its implemented policy framework serves well in the interdomain route distribution. In Junos OS, BGP is completely policy driven. The operator must explicitly configure neighbors to peer with and explicitly accept routes into BGP. Furthermore, routing policy is used to filter and modify routing information. Thus, routing policies provide complete administrative control over the routing tables.

Although, within an AS, both IGP-TE and BGP-TE provide the same set of information, BGP-TE has better scaling characteristics that are inherited from the standard BGP protocol. This makes BGP-TE a more scalable choice for acquiring multi-area/multi-AS topology information.

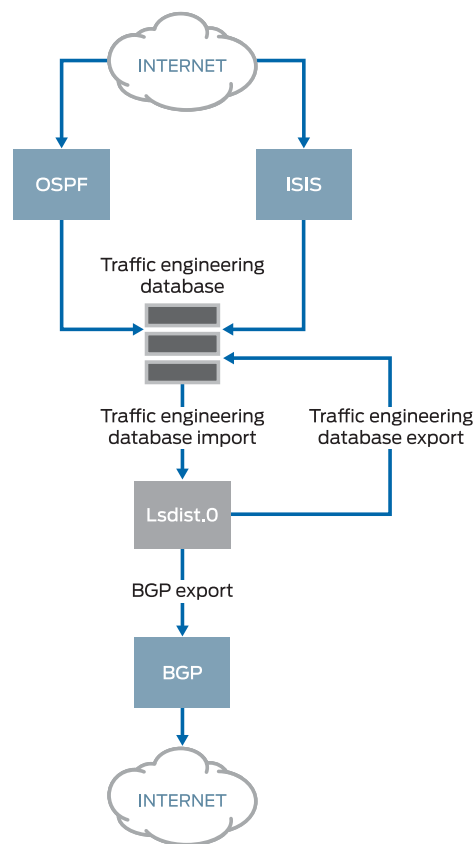
By using BGP as a solution, the IGP-acquired information is used for distribution into BGP. The ISPs can selectively expose this information with other ISPs, service providers, and content distribution networks (CDNs) through normal BGP peering. This allows for

aggregation of the IGP-acquired information across multiple areas and ASs, such that an external path computing entity can access the information by passively listening to a route reflector.

Implementation

In Junos OS, the IGPs install topology information into a database called the traffic engineering database. The traffic engineering database contains the aggregated topology information. The mechanism to distribute link-state information using BGP includes the process of advertising the traffic engineering database into BGP-TE (import), and installing entries from BGP-TE into the traffic engineering database (export).

Figure 1: Junos OS Implementation of BGP Link-State Distribution



- [Traffic Engineering Database Import on page 11](#)
- [Traffic Engineering Database Export on page 12](#)
- [Assigning Credibility Values on page 12](#)
- [BGP-TE NLRIs and TLVs on page 13](#)

Traffic Engineering Database Import

To advertise the traffic engineering database into BGP-TE, the link and node entries in the traffic engineering database are converted in the form of routes. These converted

routes are then installed by the traffic engineering database on behalf of the corresponding IGP, into a user-visible routing table called **lsdist.0**, on conditions subject to route policies. The procedure of leaking entries from the traffic engineering database into **lsdist.0** is called traffic engineering database import as illustrated in [Figure 1 on page 11](#).

There are policies to govern the traffic engineering database import process. By default, no entries are leaked from the traffic engineering database into the **lsdist.0** table.

Traffic Engineering Database Export

BGP can be configured to export or advertise routes from the **lsdist.0** table, subject to policy. This is common for any kind of route origination in BGP. In order to advertise BGP-TE into the traffic engineering database, BGP needs to be configured with the BGP-TE address family, and an export policy that selects routes for redistribution into BGP.

BGP then propagates these routes like any other NLRI. BGP peers that have the BGP-TE family configured and negotiated receive BGP-TE NRIs. BGP stores the received BGP-TE NRIs in the form of routes in the **lsdist.0** table, which is the same table that stores locally originated BGP-TE routes. The BGP-installed routes in **lsdist.0** are then distributed to other peers like any other route. Thus, the standard route selection procedure applies to BGP-TE NRIs received from multiple speakers.

To achieve interdomain TE, the routes in **lsdist.0** are leaked into the traffic engineering database through a policy. This process is called traffic engineering database export as illustrated in [Figure 1 on page 11](#).

There are policies to govern the traffic engineering database export process. By default, no entries are leaked from the **lsdist.0** table into the traffic engineering database.



NOTE: For SDN applications, such as PCE and ALTO, the BGP-TE advertised information cannot leak into the traffic engineering database of a router. In such cases, an external server that peers with the routers using BGP-TE is used to move topology information up into the sky/orchestration system that spans the network. These external servers can be deemed as BGP-TE consumers, where they receive BGP-TE routes, but do not advertise them.

Assigning Credibility Values

Once the entries are installed in the traffic engineering database, the BGP-TE learned information is made available for CSPF path computation. The traffic engineering database uses a protocol preference scheme that is based on credibility values. A protocol with a higher credibility value is preferred over a protocol with a lower credibility value. BGP-TE has the capability to advertise information learned from multiple protocols at the same time, and so in addition to the IGP-installed entries in the traffic engineering database, there can be BGP-TE installed entries that correspond to more than one protocol. The traffic engineering database export component creates a traffic engineering database protocol and credibility level for each protocol that BGP-TE supports. These credibility values are configurable in the CLI.

The credibility order for the BGP-TE protocols is as follows:

- Unknown—80
- OSPF—81
- ISIS Level 1—82
- ISIS Level 2—83
- Static—84
- Direct—85

BGP-TE NLRIs and TLVs

Like other BGP routes, BGP-TE NLRIs can also be distributed through a route reflector that speaks BGP-TE NLRI. Junos OS implements the route reflection support for the BGP-TE family.

The following is a list of supported NLRIs:

- Link NLRI
- Node NLRI
- IPv4 Prefix NLRI (receive and propagate)
- IPv6 Prefix NLRI (receive and propagate)



NOTE: Junos OS does not provide support for the route-distinguisher form of the above NLRIs.

The following is a list of supported fields in link and node NLRIs:

- Protocol-ID—NLRI originates with the following protocol values:
 - ISIS-L1
 - ISIS-L2
 - OSPF
- Identifier—This value is configurable. By default, the identifier value is set to 0.
- Local/Remote node descriptor—These include:
 - Autonomous system
 - BGP-LS Identifier—This value is configurable. By default, the BGP-LS identifier value is set to 0
 - Area-ID
 - IGP router-ID
- Link descriptors (Only for link NLRI)—This includes:

- Link Local/Remote Identifiers
- IPv4 interface address
- IPv4 neighbor address
- IPv6 neighbor/interface address—The IPv6 neighbor and interface addresses are not originated, but only stored and propagated when received.
- Multi-topology ID—This value is not originated, but stored and propagated when received.

The following is a list of supported LINK_STATE attribute TLVs:

- Link attributes:
 - Administrative group
 - Max link bandwidth
 - Max reservable bandwidth
 - Unreserved bandwidth
 - TE default metric
 - SRLG
 - The following TLVs, which are not originated, but only stored and propagated when received:
 - Opaque link attributes
 - MPLS protocol mask
 - Metric
 - Link protection type
 - Link name attribute
- Node attributes:
 - IPv4 Router-ID
 - Node flag bits—Only the overload bit is set.
 - The following TLVs, which are not originated, but only stored and propagated when received:
 - Multi-topology
 - OSPF-specific node properties
 - Opaque node properties
 - Node name

- IS-IS area identifier
- IPv6 Router-ID
- Prefix attributes—These TLVs are stored and propagated like any other unknown TLVs.

Supported and Unsupported Features

Junos OS supports the following features with link-state distribution using BGP:

- Advertisement of multiprotocol assured forwarding capability
- Transmission and reception of node and link-state BGP and BGP-TE NLRIs
- Nonstop active routing for BGP-TE NLRIs
- Policies

Junos OS does **not** support the following functionality for link-state distribution using BGP:

- Aggregated topologies, links, or nodes
- Route distinguisher support for BGP-TE NLRIs
- Multi-topology identifiers
- Multi-instance identifiers (excluding the default instance ID 0)
- Advertisement of the link and node area TLV
- Advertisement of MPLS signaling protocols
- Importing node and link information with overlapping address

Related Documentation

- [Example: Configuring Link State Distribution Using BGP on page 364](#)

PART 2

Configuring MPLS

- [MPLS Overview on page 19](#)
- [Configuring MPLS Routers on page 59](#)
- [Configuring MPLS-Signaled LSPs on page 211](#)
- [Configuring Static and Explicit-Path LSPs on page 271](#)
- [Configuring Point-to-Multipoint LSPs on page 281](#)
- [Configuring DiffServ-Aware Traffic Engineering to Achieve Service Level Guarantees on an MPLS network on page 309](#)
- [Configuring Miscellaneous MPLS Properties on page 333](#)

CHAPTER 2

MPLS Overview

- [Introduction to MPLS on page 20](#)
- [Supported MPLS Standards on page 20](#)
- [Link-Layer Support in MPLS on page 23](#)
- [MPLS and Traffic Engineering on page 24](#)
- [MPLS Label Overview on page 24](#)
- [Special MPLS Labels on page 25](#)
- [MPLS Label Allocation on page 26](#)
- [Operations on MPLS Labels on page 27](#)
- [Entropy Label Support in Mixed Mode Overview on page 28](#)
- [Routers in an LSP on page 28](#)
- [How a Packet Travels Along an LSP on page 28](#)
- [Types of LSPs on page 29](#)
- [Scope of LSPs on page 29](#)
- [Constrained-Path LSP Computation on page 29](#)
- [How CSPF Selects a Path on page 31](#)
- [CSPF Path Selection Tie-Breaking on page 32](#)
- [Computing CSPF Paths Offline on page 33](#)
- [Path Computation for LSPs on an Overloaded Router on page 33](#)
- [Computing Backup Paths for LSPs Using Fate Sharing on page 34](#)
- [Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts on page 34](#)
- [Enabling IGP Shortcuts on page 36](#)
- [LSPs Qualified in IGP Shortcut Computations on page 36](#)
- [IGP Shortcut Applications on page 36](#)
- [IGP Shortcuts and Routing Tables on page 37](#)
- [IGP Shortcuts and VPNs on page 38](#)
- [Advertising LSPs into IGP on page 38](#)
- [IP and MPLS Packets on Aggregated Interfaces on page 39](#)
- [MPLS Applications on page 40](#)

- [BGP Destinations on page 40](#)
- [IGP and BGP Destinations on page 41](#)
- [Selecting a Forwarding LSP Next Hop on page 42](#)
- [Example: Assigning Different Forwarding Next-Hop LSPs to Different Destination Prefixes on page 42](#)
- [MPLS and Routing Tables on page 43](#)
- [MPLS and Traffic Protection on page 45](#)
- [Fast Reroute Overview on page 46](#)
- [Detour Merging Process on page 48](#)
- [Detour Computations on page 49](#)
- [Fast Reroute Path Optimization on page 50](#)
- [On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview on page 50](#)

Introduction to MPLS

MPLS provides a mechanism for engineering network traffic patterns that is independent of routing tables. MPLS assigns short labels to network packets that describe how to forward them through the network. MPLS is independent of any routing protocol and can be used for unicast packets.

In the traditional Level 3 forwarding paradigm, as a packet travels from one router to the next, an independent forwarding decision is made at each hop. The IP network layer header is analyzed, and the next hop is chosen based on this analysis and on the information in the routing table. In an MPLS environment, the analysis of the packet header is performed just once, when a packet enters the MPLS cloud. The packet is then assigned to a stream, which is identified by a *label*, which is a short (20-bit), fixed-length value at the front of the packet. Labels are used as lookup indexes for the label forwarding table. For each label, this table stores forwarding information. You can associate additional information with a label—such as class-of-service (CoS) values—that can be used to prioritize packet forwarding.

Supported MPLS Standards

Junos OS substantially supports the following RFCs and Internet drafts, which define standards for MPLS and traffic engineering.

- RFC 2858, *Multiprotocol Extensions for BGP-4*
- RFC 3031, *Multiprotocol Label Switching Architecture*
- RFC 3032, *MPLS Label Stack Encoding*
- RFC 3140, *Per Hop Behavior Identification Codes*
- RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*

Only E-LSPs are supported.

- RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*
- RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol*
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

Node protection in facility backup is not supported.

- RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*
- RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*
- RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*
- RFC 4385, *Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN.*

Supported on MX Series routers with the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP.

- RFC 4875, *Extensions to RSVP-TE for Point-to-Multipoint TE LSPs*
- RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*
- RFC 5317, *Joint Working Team (JWT) Report on MPLS Architectural Considerations for a Transport Profile*
- RFC 5586, *MPLS Generic Associated Channel*
- RFC 5654, *Requirements of an MPLS Transport Profile*

The following capabilities are supported in the Junos OS implementation of MPLS Transport Profile (MPLS-TP):

- MPLS-TP OAM can send and receive packets with GAL and G-Ach, without IP encapsulation.
- Two unidirectional RSVP LSPs between a pair of routers can be associated with each other to create an associated bidirectional LSP for binding a path for the GAL and G-Ach OAM messages. A single Bidirectional Forwarding Detection (BFD) session is established for the associated bidirectional LSP.
- RFC 5712, *MPLS Traffic Engineering Soft Preemption*
- RFC 5718, *An In-Band Data Communication Network For the MPLS Transport Profile*
- RFC 5860, *Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks*
- RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*
- RFC 5921, *A Framework for MPLS in Transport Networks*
- RFC 5950, *Network Management Framework for MPLS-based Transport Networks*
- RFC 5951, *Network Management Requirements for MPLS-based Transport Networks*
- RFC 5960, *MPLS Transport Profile Data Plane Architecture*

- RFC 6215, *MPLS Transport Profile User-to-Network and Network-to-Network Interfaces*
- RFC 6291, *Guidelines for the Use of the “OAM” Acronym in the IETF.*
- RFC 6370, *MPLS Transport Profile (MPLS-TP) Identifiers*
- RFC 6371, *Operations, Administration, and Maintenance Framework for MPLS-Based Transport Networks.*
- RFC 6372, *MPLS Transport Profile (MPLS-TP) Survivability Framework*
- RFC 6373, *MPLS-TP Control Plane Framework*
- RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*

Only Point-to-Multipoint LSPs are supported.

- RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*
- RFC 6425, *Detecting Data-Plane Failures in Point-to-Multipoint MPLS - Extensions to LSP Ping*
- RFC 6426, *MPLS On-Demand Connectivity Verification and Route Tracing*
- RFC 6428, *Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for the MPLS Transport Profile*
- Internet draft draft-ietf-mpls-rsvp-te-no-php-oob-mapping-01.txt, *Non PHP behavior and Out-of-Band Mapping for RSVP-TE LSPs*

The following RFCs and Internet drafts do not define standards, but provide information about MPLS, traffic engineering, and related technologies. The IETF classifies them variously as “Experimental,” “Historic,” or “Informational.”

- RFC 2547, *BGP/MPLS VPNs*
 - RFC 2702, *Requirements for Traffic Engineering Over MPLS*
 - RFC 2917, *A Core MPLS IP VPN Architecture*
 - RFC 3063, *MPLS Loop Prevention Mechanism*
 - RFC 3208, *PGM Reliable Transport Protocol Specification*
- Only the network element is supported.
- RFC 3469, *Framework for Multi-Protocol Label Switching (MPLS)-based Recovery*
 - RFC 3564, *Requirements for Support of Differentiated Services-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*
 - Internet draft draft-martini-l2circuit-encap-mpls-11.txt, *Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks*

Junos OS differs from the Internet draft in the following ways:

- A packet with a sequence number of 0 is treated as out of sequence.
- Any packet that does not have the next incremental sequence number is considered out of sequence.
- When out-of-sequence packets arrive, the expected sequence number for the neighbor is set to the sequence number in the Layer 2 circuit control word.
- Internet draft draft-martini-l2circuit-trans-mpls-19.txt, *Transport of Layer 2 Frames Over MPLS*
- Internet draft draft-raggarwa-mpls-p2mp-te-02.txt, *Establishing Point to Multipoint MPLS TE LSPs*

The features discussed in the indicated sections of the draft are not supported:

- Nonadjacent signaling for branch LSPs (section 7.1)
- Make-before-break and fast reroute (section 9)
- LSP hierarchy using point-to-point LSPs (section 10)

Related Documentation

- [Supported GMPLS Standards on page 673](#)
- [Supported LDP Standards on page 520](#)
- [Supported RSVP Standards on page 460](#)
- [Accessing Standards Documents on the Internet](#)

Link-Layer Support in MPLS

MPLS supports the following link-layer protocols, which are all supported in the Junos OS MPLS implementation:

- Point-to-Point Protocol (PPP)—Protocol ID 0x0281, Network Control Protocol (NCP) protocol ID 0x8281.
- Ethernet/Cisco High-level Data Link Control (HDLC)—Ethernet type 0x8847.
- Asynchronous Transfer Mode (ATM)—Subnetwork attachment point encoded (SNAP-encoded) Ethernet type 0x8847. Support is included for both point-to-point mode or nonbroadcast multiaccess (NBMA) mode. Support is not included for encoding MPLS labels as part of ATM virtual path identifier/virtual circuit identifier (VPI/VCI).
- Frame Relay—SNAP-encoded, Ethernet type 0x8847. Support is not included for encoding MPLS labels as part of Frame Relay data-link connection identifier (DLCI).
- Generic routing encapsulation (GRE) tunnel—Ethernet type 0x8847.

MPLS and Traffic Engineering

Traffic engineering allows you to control the path that data packets follow, bypassing the standard routing model, which uses routing tables. Traffic engineering moves flows from congested links to alternate links that would not be selected by the automatically computed destination-based shortest path. With traffic engineering, you can:

- Make more efficient use of expensive long-haul fibers.
- Control how traffic is rerouted in the face of single or multiple failures.
- Classify critical and regular traffic on a per-path basis.

The core of the traffic engineering design is based on building label-switched paths (LSPs) among routers. An LSP is connection-oriented, like a virtual circuit in Frame Relay or ATM. LSPs are not reliable: Packets entering an LSP do not have delivery guarantees, although preferential treatment is possible. LSPs also are similar to unidirectional tunnels in that packets entering a path are encapsulated in an envelope and switched across the entire path without being touched by intermediate nodes. LSPs provide fine-grained control over how packets are forwarded in a network. To provide reliability, an LSP can use a set of primary and secondary paths.

LSPs can be configured for BGP traffic only (traffic whose destination is outside of an autonomous system [AS]). In this case, traffic within the AS is not affected by the presence of LSPs. LSPs can also be configured for both BGP and interior gateway protocol (IGP) traffic; therefore, both intra-AS and inter-AS traffic is affected by the LSPs.

This section discusses the following topics:

- [MPLS Label Overview on page 24](#)
- [MPLS Label Allocation on page 26](#)
- [Routers in an LSP on page 28](#)
- [How a Packet Travels Along an LSP on page 28](#)
- [Types of LSPs on page 29](#)
- [Scope of LSPs on page 29](#)
- [Constrained-Path LSP Computation on page 29](#)
- [Path Computation for LSPs on an Overloaded Router on page 33](#)
- [Computing Backup Paths for LSPs Using Fate Sharing on page 34](#)
- [Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts on page 34](#)
- [Advertising LSPs into IGP on page 38](#)

MPLS Label Overview

Packets traveling along an LSP are identified by a label—a 20-bit, unsigned integer in the range 0 through 1,048,575. For push labels on ingress routers, no labels in this range are

restricted. For incoming labels on the transit static LSP, the label value is restricted to 1,000,000 through 1,048,575.

On MX Series, PTX Series, and T Series routers, the value for entropy and flow labels is restricted to 16 through 1,048,575.

Special MPLS Labels

Some of the reserved labels (in the 0 through 15 range) have well-defined meanings. For more complete details, see RFC 3032, *MPLS Label Stack Encoding*.

- 0, IPv4 Explicit Null label—This value is legal only when it is the sole label entry (no label stacking). It indicates that the label must be popped upon receipt. Forwarding continues based on the IP version 4 (IPv4) packet.
- 1, Router Alert label—When a packet is received with a top label value of 1, it is delivered to the local software module for processing.
- 2, IPv6 Explicit Null label—This value is legal only when it is the sole label entry (no label stacking). It indicates that the label must be popped on receipt. Forwarding continues based on the IP version 6 (IPv6) packet.
- 3, Implicit Null label—This label is used in the control protocol (LDP or RSVP) only to request label popping by the downstream router. It never actually appears in the encapsulation. Labels with a value of 3 should not be used in the data packet as real labels. No payload type (IPv4 or IPv6) is implied with this label.
- 4 through 6—Unassigned.
- 7, Entropy label indicator—This label is used when an Entropy label is in the label stack and precedes the Entropy label.
- 8 through 15—Unassigned.

Special labels are commonly used between the egress and penultimate routers of an LSP. If the LSP is configured to carry IPv4 packets only, the egress router might signal the penultimate router to use 0 as a final-hop label. If the LSP is configured to carry IPv6 packets only, the egress router might signal the penultimate router to use 2 as a final-hop label.

The egress router might simply signal the penultimate router to use 3 as the final label, which is a request to perform penultimate-hop label popping. The egress router will not process a labeled packet; rather, it receives the payload (IPv4, IPv6, or others) directly, reducing one MPLS lookup at egress.

For label-stacked packets, the egress router receives an MPLS label packet with its top label already popped by the penultimate router. The egress router cannot receive label-stacked packets that use label 0 or 2. It typically requests label 3 from the penultimate router.

MPLS Label Allocation

In the Junos OS, label values are allocated per router or switch—the rest of this explanation uses router to cover both. The display output shows only the label (for example, **01024**). Labels for multicast packets are independent of those for unicast packets. Currently, the Junos OS does not support multicast labels.

Labels are assigned by downstream routers relative to the flow of packets. A router receiving labeled packets (the next-hop router) is responsible for assigning incoming labels. A received packet containing a label that is unrecognized (unassigned) is dropped. For unrecognized labels, the router does not attempt to unwrap the label to analyze the network layer header, nor does it generate an Internet Control Message Protocol (ICMP) destination unreachable message.

A packet can carry a number of labels, organized as a last-in, first-out stack. This is referred to as a *label stack*. At a particular router, the decision about how to forward a labeled packet is based exclusively on the label at the top of the stack.

Figure 2 on page 26 shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.

Figure 2: Label Encoding





Figure 3 on page 27 illustrates the purpose of the class-of-service bits (also known as the EXP or experimental bits). Bits 20 and 21 specify the queue number. Bit 22 is the packet loss priority (PLP) bit used to specify the random early detection (RED) drop profile. For more information about class of service and the class-of-service bits, see “Configuring Class of Service for MPLS LSPs” on page 244.

Figure 3: Class-of-Service Bits



Related Documentation

- [per-prefix-label on page 921](#)

Operations on MPLS Labels

The router supports the following label operations:

- **Push**—Add a new label to the top of the packet. For IPv4 packets, the new label is the first label. The time-to-live (TTL) and s bits are derived from the IP packet header. The MPLS class of service (CoS) is derived from the queue number. If the push operation is performed on an existing MPLS packet, you will have a packet with two or more labels. This is called label stacking. The top label must have its s bit set to 0, and might derive CoS and TTL from lower levels. The new top label in a label stack always initializes its TTL to 255, regardless of the TTL value of lower labels.
- **Pop**—Remove the label from the beginning of the packet. Once the label is removed, the TTL is copied from the label into the IP packet header, and the underlying IP packet is forwarded as a native IP packet. In the case of multiple labels in a packet (label stacking), removal of the top label yields another MPLS packet. The new top label might derive CoS and TTL from a previous top label. The popped TTL value from the previous top label is not written back to the new top label.
- **Swap**—Replace the label at the top of the label stack with a new label. The S and CoS bits are copied from the previous label, and the TTL value is copied and decremented (unless the **no-decrement-ttl** or **no-propagate-ttl** statement is configured). A transit router supports a label stack of any depth.
- **Multiple Push**—Add multiple labels (up to three) on top of existing packets. This operation is equivalent to pushing multiple times.
- **Swap and Push**—Replace the existing top of the label stack with a new label, and then push another new label on top.

Entropy Label Support in Mixed Mode Overview

The entropy label helps transit routers load-balance MPLS traffic across ECMP paths or link aggregation groups. The entropy label introduces a load-balancing label to be used by routers to load balance traffic rather than relying on deep packet inspection, reducing the packet processing requirements in the forwarding plane at the expense of increased label stack depth. Junos OS supports the entropy label only for MX Series routers with MPCs or MICs and can be enabled with enhanced-ip mode. But, this leads to a packet drop if the core-facing interface has an entropy label configured on the MPC or MIC and the other end of this core-facing connection has a DPC line card. In order to avoid this, the entropy label is now supported in mixed mode where the entropy label can be configured without enhanced-ip configuration. This allows MX Series router DPCs to support a pop out entropy label. However, this does not support a flow label.

Routers in an LSP

Each router in an LSP performs one of the following functions:

- Ingress router—The router at the beginning of an LSP. This router encapsulates IP packets with an MPLS Layer 2 frame and forwards it to the next router in the path. Each LSP can have only one ingress router.
- Egress router—The router at the end of an LSP. This router removes the MPLS encapsulation, thus transforming it from an MPLS packet to an IP packet, and forwards the packet to its final destination using information in the IP forwarding table. Each LSP can have only one egress router. The ingress and egress routers in an LSP cannot be the same router.
- Transit router—Any intermediate router in the LSP between the ingress and egress routers. A transit router forwards received MPLS packets to the next router in the MPLS path. An LSP can contain zero or more transit routers, up to a maximum of 253 transit routers in a single LSP.

A single router can be part of multiple LSPs. It can be the ingress or egress router for one or more LSPs, and it also can be a transit router in one or more LSPs. The functions that each router supports depend on your network design.

How a Packet Travels Along an LSP

When an IP packet enters an LSP, the ingress router examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label.

The packet is then forwarded to the next router in the LSP. This router and all subsequent routers in the LSP do not examine any of the IP routing information in the labeled packet. Rather, they use the label to look up information in their label forwarding table. They then replace the old label with a new label and forward the packet to the next router in the path.

When the packet reaches the egress router, the label is removed, and the packet again becomes a native IP packet and is again forwarded based on its IP routing information.

Types of LSPs

There are three types of LSPs:

- Static LSPs—For static paths, you must manually assign labels on all routers involved (ingress, transit, and egress). No signaling protocol is needed. This procedure is similar to configuring static routes on individual routers. Like static routes, there is no error reporting, liveliness detection, or statistics reporting.
- LDP-signaled LSPs—See [“LDP Introduction” on page 519](#).
- RSVP-signaled LSPs—For signaled paths, RSVP is used to set up the path and dynamically assign labels. (RSVP signaling messages are used to set up signaled paths.) You configure only the ingress router. The transit and egress routers accept signaling information from the ingress router, and they set up and maintain the LSP cooperatively. Any errors encountered while establishing an LSP are reported to the ingress router for diagnostics. For signaled LSPs to work, a version of RSVP that supports tunnel extensions must be enabled on all routers.

There are two types of RSVP-signaled LSPs:

- Explicit-path LSPs—All intermediate hops of the LSP are manually configured. The intermediate hops can be strict, loose, or any combination of the two. Explicit path LSPs provide you with complete control over how the path is set up. They are similar to static LSPs but require much less configuration.
- Constrained-path LSPs—The intermediate hops of the LSP are automatically computed by the software. The computation takes into account information provided by the topology information from the IS-IS or OSPF link-state routing protocol, the current network resource utilization determined by RSVP, and the resource requirements and constraints of the LSP. For signaled constrained-path LSPs to work, either the IS-IS or OSPF protocol and the IS-IS or OSPF traffic engineering extensions must be enabled on all routers.

Scope of LSPs

For constrained-path LSPs, the LSP computation is confined to one IGP domain, and cannot cross any AS boundary. This prevents an AS from extending its IGP into another AS.

Explicit-path LSPs, however, can cross as many AS boundaries as necessary. Because intermediate hops are manually specified, the LSP does not depend on the IGP topology or a local forwarding table.

Constrained-Path LSP Computation

The Constrained Shortest Path First (CSPF) algorithm is an advanced form of the shortest-path-first (SPF) algorithm used in OSPF and IS-IS route computations. CSPF

is used in computing paths for LSPs that are subject to multiple constraints. When computing paths for LSPs, CSPF considers not only the topology of the network, but also the attributes of the LSP and the links, and it attempts to minimize congestion by intelligently balancing the network load.

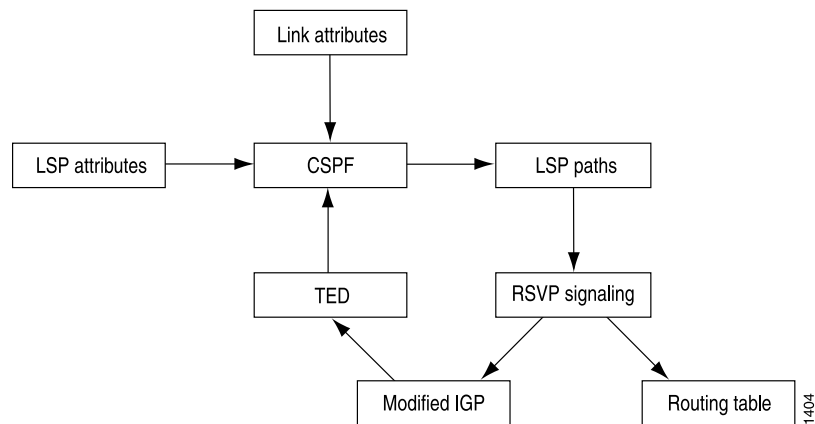
The constraints that CSPF considers include:

- LSP attributes
 - Administrative groups (that is, link color requirements)
 - Bandwidth requirements
 - Explicit route (strict or loose)
 - Hop limitations
 - Priority (setup and hold)
- Link attributes
 - Administrative groups (that is, link colors assigned to the link)
 - Reservable bandwidth of the links (static bandwidth minus the currently reserved bandwidth)

The data that CSPF considers comes from the following sources:

- Traffic engineering database—Provides CSPF with up-to-date topology information, the current reservable bandwidth of links, and the link colors. For the CSPF algorithm to perform its computations, a link-state IGP (such as OSPF or IS-IS) with special extensions is needed. For CSPF to be effective, the link-state IGP on all routers must support the special extensions. While building the topology database, the extended IGP must take into consideration the current LSPs and must flood the route information everywhere. Because changes in the reserved link bandwidth and link color cause database updates, an extended IGP tends to flood more frequently than a normal IGP. See [Figure 4 on page 31](#) for a diagram of the relationships between these components.
- Currently active LSPs—Includes all the LSPs that should originate from the router and their current operational status (up, down, or timeout).

Figure 4: CSPF Computation Process



This section discusses the following topics:

- [How CSPF Selects a Path on page 31](#)
- [CSPF Path Selection Tie-Breaking on page 32](#)
- [Computing CSPF Paths Offline on page 33](#)

How CSPF Selects a Path

To select a path, CSPF follows certain rules. The rules are as follows:

1. Computes LSPs one at a time, beginning with the highest priority LSP (the one with the lowest setup priority value). Among LSPs of equal priority, CSPF services the LSPs in alphabetical order of the LSP names.
2. Prunes the traffic engineering database of all the links that are not full duplex and do not have sufficient reservable bandwidth.
3. If the LSP configuration includes the **include** statement, prunes all links that do not share any included colors.
4. If the LSP configuration includes the **exclude** statement, prunes all links that contain excluded colors. If the link does not have a color, it is accepted.
5. If several paths have equal cost, chooses the one whose last-hop address is the same as the LSP's destination.
6. If several equal cost paths remain, selects the one with the fewest number of hops.
7. If several equal-cost paths remain, applies the CSPF load-balancing rule configured on the LSP (least fill, most fill, or random).

CSPF finds the shortest path toward the LSP's egress router, taking into account explicit-path constraints. For example, if the path must pass through Router A, two separate SPF computations are performed, one from the ingress router to Router A, the other from Router A to the egress router. All CSPF rules are applied to both computations.

- Related Documentation**
- [Configuring CSPF Tie Breaking on page 232](#)
 - [CSPF Path Selection Tie-Breaking on page 32](#)

CSPF Path Selection Tie-Breaking

If more than one path is still available after the CSPF rules ([“How CSPF Selects a Path” on page 31](#)) have been applied, a tie-breaking rule is applied to choose the path for the LSP. The rule used depends on the configuration. There are three tie-breaking rules:

- **Random**—One of the remaining paths is picked at random. This rule tends to place an equal number of LSPs on each link, regardless of the available bandwidth ratio. This is the default behavior.
- **Least fill**—The path with the largest minimum available bandwidth ratio is preferred. This rule tries to equalize the reservation on each link.
- **Most fill**—The path with the smallest minimum available bandwidth ratio is preferred. This rule tries to fill a link before moving on to alternative links.

The following definitions describe how a figure for minimum available bandwidth ratio is derived for the least fill and most fill rules:

- Reservable bandwidth = bandwidth of link x subscription factor of link
- Available bandwidth = reservable bandwidth – (sum of the bandwidths of the LSPs traversing the link)
- Available bandwidth ratio = available bandwidth/reservable bandwidth
- Minimum available bandwidth ratio (for a path) = the smallest available bandwidth ratio of the links in a path



NOTE: For the least fill or most fill behaviors to be used, the paths must have their bandwidth (specified using the `bandwidth` statement at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level) or minimum bandwidth (specified using the `minimum-bandwidth` statement at the `[edit protocols mpls label-switched-path lsp-name auto-bandwidth]` hierarchy level) configured to a value greater than 0. If the bandwidth or minimum bandwidth for the paths is either not configured or configured as 0, the minimum available bandwidth cannot be calculated and the random path selection behavior is used instead.

- Related Documentation**
- [How CSPF Selects a Path on page 31](#)
 - [Configuring CSPF Tie Breaking on page 232](#)
 - [Configuring the Bandwidth Value for LSPs on page 256](#)
 - [Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 259](#)

Computing CSPF Paths Offline

The Junos OS provides online, real-time CSPF computation only; each router performs CSPF calculations independent of the other routers in the network. These calculations are based on currently available topology information—information that is usually recent, but not completely accurate. LSP placements are locally optimized, based on current network status.

To optimize links globally across the network, you can use an offline tool to perform the CSPF calculations and determine the paths for the LSPs. You can create such a tool yourself, or you can modify an existing network design tool to perform these calculations. You should run the tool periodically (daily or weekly) and download the results into the router. An offline tool should take the following into account when performing the optimized calculations:

- All the LSP's requirements
- All link attributes
- Complete network topology

Path Computation for LSPs on an Overloaded Router

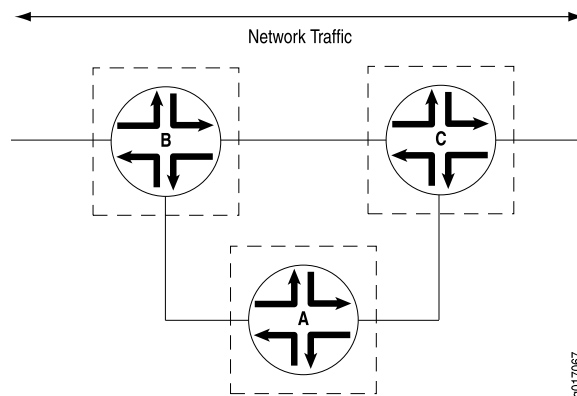
An overloaded router is a router running IS-IS with its overload bit set in its IS-IS configuration. In this case, an MPLS LSP specifically refers to an RSVP-signaled or LDP-signaled LSP. In the case of RSVP, it applies to both CSPF and non-CSPF LSPs.

You cannot establish transit LSPs through an overloaded router. However, you can configure ingress and egress LSPs through an overloaded router.



NOTE: When you set the overload bit on an IS-IS router, all LSPs transiting through it are recomputed and rerouted away from it. If the recomputation fails, no additional attempt to reconfigure the LSP is made, and the affected LSPs are disconnected.

An example of when you might want to establish transit LSPs through an overloaded router is illustrated in [Figure 5 on page 34](#), which shows an aggregation router (Router A) dual-homed on two core routers (Router B and Router C). You want to include the aggregation router in the LSP mesh, but transit LSPs should not pass through it, because it is a less capable router with relatively low-bandwidth uplinks to the core. Certain failure and rerouting scenarios could make it impossible for the aggregation router to establish some of its LSPs. Consequently, you run the router in a steady state with the overload bit set, but you are still able to establish ingress and egress LSPs through it.

Figure 5: Aggregation Router A Dual-Homed on Core Routers B and C

Computing Backup Paths for LSPs Using Fate Sharing

Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group.

Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination.

For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously.


- Related Documentation**
- [Configuring the Ingress Router for MPLS-Signaled LSPs on page 60](#)
 - [fate-sharing on page 871](#)

Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts

Link-state protocols, such as OSPF and IS-IS, use the shortest-path-first (SPF) algorithm to compute the shortest-path tree to all nodes in the network. The results of such computations can be represented by the destination node, next-hop address, and output interface, where the output interface is a physical interface. Label-switched paths (LSPs) can be used to augment the SPF algorithm, for the purposes of resolving BGP next hops. On the node performing the calculations, LSPs appear to be logical interfaces directly connected to remote nodes in the network. If you configure the interior gateway protocol (IGP) to treat LSPs the same as a physical interface and use the LSPs as a potential output interface, the SPF computation results are represented by the destination node and output LSP, effectively using the LSP as a shortcut through the network to the destination.


As an illustration, begin with a typical SPF tree (see [Figure 6 on page 35](#)).

Figure 6: Typical SPF Tree, Sourced from Router A



If an LSP connects Router A to Router D and if IGP shortcuts are enabled on Router A, you might have the SPF tree shown in [Figure 7 on page 35](#).

Figure 7: Modified SPF Tree, Using LSP A–D as a Shortcut



Note that Router D is now reachable through LSP A–D.

When computing the shortest path to reach Router D, Router A has two choices:


- Use IGP path A–B–D.
- Use LSP A–D.

Router A decides between the two choices by comparing the IGP metrics for path A–B–D with the LSP metrics for LSP A–D. If the IGP metric is lower, path A–B–D is chosen ([Figure 6 on page 35](#)). If the LSP metric is lower, LSP A–D is used ([Figure 7 on page 35](#)). If both metrics are equal, LSP A–D is chosen because LSPs are preferred over IGP paths.

Note that Routers E and F are also reachable through LSP A–D, because they are downstream from Router D in the SPF tree.

Assuming that another LSP connects Router A to Router E, you might have the SPF tree shown in [Figure 8 on page 36](#).

Figure 8: Modified SPF Tree, Using LSP A–D and LSP A–E as Shortcuts



- Related Documentation**
- [traffic-engineering](#)
 - [Understanding OSPF Support for Traffic Engineering](#)

Enabling IGP Shortcuts

IGP shortcuts are supported for both IS-IS and OSPF. A link-state protocol is required for IGP shortcuts. Shortcuts are disabled by default. You can enable IGP shortcuts on a per-router basis; you do not need to enable shortcuts globally. A router's shortcut computation does not depend on another router performing similar computations, and shortcuts performed by other routers are irrelevant.

- Related Documentation**
- [Example: Enabling IS-IS Traffic Engineering Support](#)
 - [Example: Enabling OSPF Traffic Engineering Support](#)
 - [Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts on page 34](#)

LSPs Qualified in IGP Shortcut Computations

Not all LSPs are used in IGP shortcuts. Only those LSPs whose egress point (using the **to** statement) matches the router ID of the egress node are considered. Other LSPs, whose egress point matches the egress node interface address, are ignored in IGP shortcuts.

There are exceptions, however. If an LSP has an alias egress point (using the **install** statement) and it matches certain router IDs, it is included in the shortcut computation as well. If multiple equal metric LSPs destined to the same router ID exist, traffic can load-share among them.


IGP Shortcut Applications

You can use shortcuts to engineer traffic traveling toward destination nodes that do not support MPLS LSPs. For example, in [Figure 8 on page 36](#), traffic traveling toward Router F enters LSP A–E. You can control traffic between Router A and Router F by manipulating LSP A–E; you do not need to explicitly set up an LSP between Router A and Router F.

In [Figure 9 on page 37](#), all traffic from Region 1 to Region 2 traverses LSP A–B if IGP shortcuts are enabled on the ingress router (Router A), permitting aggregation of


interregional traffic into one LSP. To perform traffic engineering on the interregional traffic, you have to manipulate LSP A-B only, which avoids creating n^2 LSPs from all routers in Region 1 to all routers in Region 2 and allows efficient resource controls on the backbone network.

Figure 9: IGP Shortcuts



Shortcuts allow you to deploy LSPs into a network in an incremental, hierarchical fashion. In [Figure 10 on page 37](#), each region can choose to implement traffic engineering LSPs independently, without requiring cooperation from other regions. Each region can choose to deploy intraregion LSPs to fit the region's bandwidth needs, at the pace appropriate for the region.

Figure 10: IGP Shortcuts in a Bigger Network



When intraregion LSPs are in place, interregional traffic automatically traverses the intraregion LSPs as needed, eliminating the need for a full mesh of LSPs between edge routers. For example, traffic from Router A to Router D traverses LSPs A-B, B-C, and C-D.

IGP Shortcuts and Routing Tables

IGP typically performs two independent computations. The first is performed without considering any LSP. The result of the computation is stored in the `inet.0` table. This step is no different from traditional SPF computations and is always performed even if IGP shortcut is disabled.

The second computation is performed considering only LSPs as a logical interface. Each LSP's egress router is considered. The list of destinations whose shortest path traverses the egress router (established during the first computation) is placed in the `inet.3` routing table. These destinations are given the egress router of the LSP as a next hop, enabling BGP on the local router to use these LSPs to access BGP next hops beyond the egress router. Normally, BGP can use only LSPs that terminate at the BGP next hop. Note that BGP is the only protocol that uses the `inet.3` routing table. Other protocols will not route traffic through these LSPs.

If traffic engineering for IGP and BGP is enabled (see [“IGP and BGP Destinations” on page 41](#)), IGP moves all routes in `inet.3` into `inet.0`, merging all routes while emptying the `inet.3` table. The number of routes in `inet.0` will be exactly the same as before. Route next-hops can traverse a physical interface, an LSP, or the combination of the two if the metrics are equal.

IGP shortcuts are enabled on a per-node basis. You do not need to coordinate with other nodes.

IGP Shortcuts and VPNs

You can configure IGP shortcuts for either IS-IS or OSPF. IGP shortcuts allow the IGP to use an LSP as the next hop instead of the IGP route. IGP shortcuts can also be enabled for VPNs by also specifying the **bgp-igp-both-ribs** or **mpls-forwarding** options for the **traffic-engineering** statement at the **[edit protocols mpls]** hierarchy level. VPNs are dependant on routes stored in the inet.3 routing table. The **bgp-igp** option for the **traffic-engineering** statement moves all routes from the inet.3 routing table to the inet.0 routing table and is therefore incompatible with VPNs.


- Related Documentation**
- [Configuring Traffic Engineering for LSPs on page 336](#)
 - *traffic-engineering*
 - *Understanding OSPF Support for Traffic Engineering*

Advertising LSPs into IGPs

You can configure your IGP to treat an LSP as a link. IGP shortcuts allow only the ingress router of an LSP to use the LSP in its SPF computation. However, other routers on the network do not know of the existence of that LSP, so they cannot use it. This can lead to suboptimal traffic engineering. In addition, only BGP can use an IGP shortcut to an LSP. When you advertise an LSP as a link into the IGP, all traffic can traverse it, and all routers know about it.

As an example, consider the network shown in [Figure 11 on page 38](#).

Figure 11: SPF Computations with Advertised LSPs



Assume that Router A is computing a path to Router D. The link between Router E and Router F has a metric of 20; all other links have a metric of 10. Here, the path chosen by Router A is A–B–C–D, which has a metric of 30, instead of A–E–F–D, which has a metric of 40.

If Router E has an LSP to Router D with a metric of 15, you want traffic from Router A to Router D to use the path A–E–D, which has a metric of 25, instead of the path A–B–C–D. However, because Router A does not know about the LSP between Router E and Router D, it cannot route traffic through this path.

For all routers on the network to know about the LSP between Router E and Router D, you need to advertise it. This advertisement announces the LSP as a unidirectional, point-to-point link in the link-state database, and all routers can compute paths using the LSP. The link-state database maintains information about the AS topology and contains information about the router's local state (for example, the router's usable interfaces and reachable neighbors). In [Figure 11 on page 38](#), Router A will see the link from Router E to Router D and route traffic along this lower-metric path.

Because an LSP is announced as a unidirectional link, you might need to configure a reverse LSP (one that starts at the egress router and ends at the ingress router) so that the SPF bidirectionality check succeeds. As a step in the SPF computation, IS-IS considers a link from Router E to Router D. Before IS-IS uses any link, it verifies that there is a link from Router D to Router E (there is bidirectional connectivity between router E and D). Otherwise, the SPF computation will not use an announced LSP.

When an LSP is advertised to the IGP, the advertising router uses the LSP as the forwarding path for regular routes after installing them in the inet.0 routing table. All packets traversing the router could be forwarded through the LSP. Conversely, IGP shortcuts are used only to forward packets that are following BGP routes.



NOTE: Do not configure IGP shortcuts and advertise LSPs to the IGP at the same time.

IP and MPLS Packets on Aggregated Interfaces

You can send IP and MPLS packets over aggregated interfaces. To the IP or MPLS session, there is a single LSP composed of the aggregated interfaces. Packets sent to an LSP that is part of an aggregated interface are redistributed over the aggregated member interfaces.

Sending IP and MPLS packets over aggregated interfaces has the following benefits:

- **Bandwidth aggregation**—You can increase the number of MPLS packet flows sent over each connection. In MPLS, a set of packets sharing the same label is considered a part of the same flow.
- **Link redundancy**—If a link or a line card failure affects an aggregate member link, the traffic flowing across that link is immediately forwarded across one of the remaining links.

The Junos OS supports aggregated SONET and Ethernet interfaces.

Note that the Junos implementation of IP and MPLS over aggregated interfaces (aggregated Ethernet devices only) complies with IEEE 802.3ad.

For information about how to configure aggregated Ethernet or aggregated SONET interfaces, see *Ethernet Interfaces Feature Guide for Routing Devices* and *Configuring Aggregated SONET/SDH Interfaces*.

- Related Documentation**
- *Ethernet Interfaces Feature Guide for Routing Devices*
 - *Configuring Aggregated SONET/SDH Interfaces*

MPLS Applications

In the Junos OS implementation of MPLS, establishing an LSP installs on the ingress router a host route (a 32-bit mask) toward the egress router. The address of the host route is the destination address of the LSP. By default, the route has a preference value of 7, a value that is higher than all routes except direct interface and static routes. The 32-bit mask ensures that the route is more specific (that is, a longer match) than all other subnet routes. The host routes can be used to traffic-engineer BGP destinations only, or both IGP and BGP destinations.

This section discusses the following topics:

- [BGP Destinations on page 40](#)
- [IGP and BGP Destinations on page 41](#)
- [Selecting a Forwarding LSP Next Hop on page 42](#)

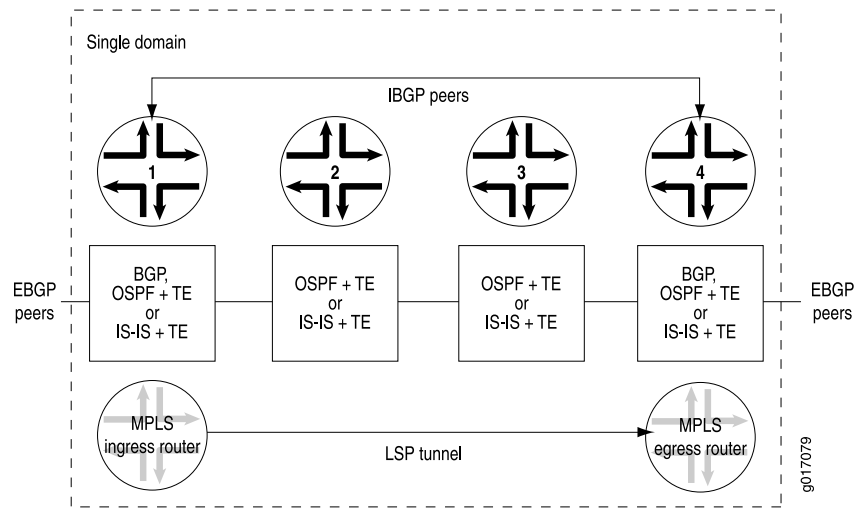
BGP Destinations

You can configure MPLS to control the paths that traffic takes to destinations outside an AS.

Both IBGP and EBGP take advantage of the LSP host routes without requiring extra configuration. BGP compares the BGP next-hop address with the LSP host route. If a match is found, the packets for the BGP route are label-switched over the LSP. If multiple BGP routes share the same next-hop address, all the BGP routes are mapped to the same LSP route, regardless of which BGP peer the routes are learned from. If the BGP next-hop address does not match an LSP host route, BGP routes continue to be forwarded based on the IGP routes within the routing domain. In general, when both an LSP route and an IGP route exist for the same BGP next-hop address, the one with the lowest preference is chosen.

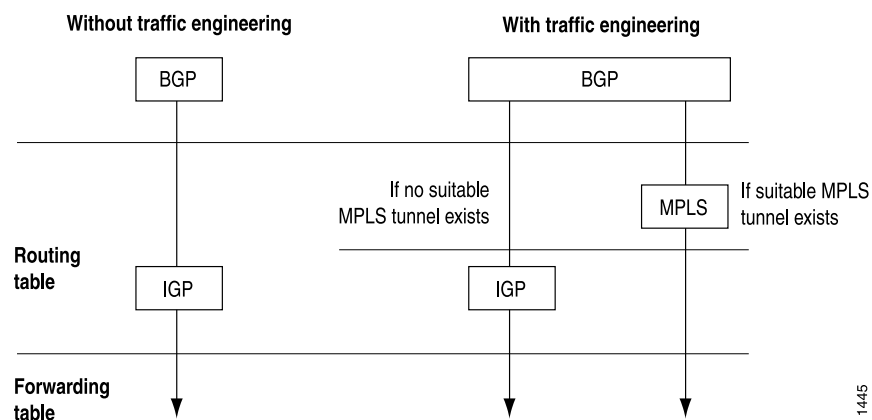
[Figure 12 on page 41](#) shows an MPLS topology that illustrates how MPLS and LSPs work. This topology consists of a single domain with four routers. The two routers at the edges of the domain, Router 1 and Router 4, are running EBGP to communicate with peers outside the domain and IBGP to communicate between themselves. For intradomain communication, all four routers are running an IGP. Finally, an LSP tunnel exists from Router 1 to Router 4.

Figure 12: MPLS Application Topology



When BGP on Router 1 receives prefixes from Router 4, it must determine how to reach a BGP next-hop address. Typically, when traffic engineering is not enabled, BGP uses IGP routes to determine how to reach next-hop addresses. (See the left side of [Figure 13 on page 41](#).) However, when traffic engineering is enabled, if the BGP next-hop matches the LSP tunnel endpoint (that is, the MPLS egress router), those prefixes enter the LSP tunnel. (To track these prefixes, look at the **Active Route** field in the `show mpls lsp` command output or at the output of the `show route label-switched-path path-name` command.) If the BGP next hop does not match an LSP tunnel endpoint, those prefixes are sent following the IGP's shortest path. (See [Figure 13 on page 41](#).)

Figure 13: How BGP Determines How to Reach Next-Hop Addresses



IGP and BGP Destinations

You can configure MPLS to control the paths that traffic takes to destinations within an AS.

When traffic engineering is for BGP destinations only, the MPLS host routes are installed in the inet.3 routing table (see [Figure 14 on page 44](#)), separate from the routes learned from other routing protocols. Not all inet.3 routes are downloaded into the forwarding

table. Packets directly addressed to the egress router do not follow the LSP, which prevents routes learned from LSPs from overriding routes learned from IGP or other sources.

Traffic within a domain, including BGP control traffic between BGP peers, is not affected by LSPs. MPLS affects interdomain traffic only; that is, it affects only those BGP prefixes that are learned from an external domain. MPLS does not disrupt intradomain traffic, so IS-IS or OSPF routes remain undisturbed. If you issue a **ping** or **traceroute** command to any destination within the domain, the **ping** or **traceroute** packets follow the IGP path. However, if you issue a **ping** or **traceroute** command from Router 1 in [Figure 12 on page 41](#) (the LSP ingress router) to a destination outside of the domain, the packets use the LSP tunnel.

When traffic engineering for IGP and BGP destinations is enabled, the MPLS host routes are installed in the inet.0 table (see [Figure 15 on page 45](#)) and downloaded into the forwarding table. Any traffic destined to the egress router could enter the LSP. In effect, it moves all the routes in inet.3 into inet.0, causing the inet.3 table to be emptied.

RSVP packets automatically avoid all MPLS LSPs, including those established by RSVP or LDP. This prevents placing one RSVP session into another LSP, or in other words, nesting one LSP into another.

Selecting a Forwarding LSP Next Hop

If more than one LSP tunnel to a BGP next hop exists, the prefixes learned from the BGP next hop are randomly divided among the LSP tunnels. To control which LSP BGP uses to forward data for a given prefix, use the **install-nexthop** statement in the export policy applied to the forwarding table.

- Related Documentation**
- *Configuring Policies for Layer 2 Circuits*
 - *install-nexthop*

Example: Assigning Different Forwarding Next-Hop LSPs to Different Destination Prefixes

Assign different forwarding next-hop LSPs to different destination prefixes learned from BGP.

```
routing-options {
  router-id 10.10.20.101;
  autonomous-system 2;
  forwarding-table {
    export forwarding-policy;
  }
}
policy-options {
  policy-statement forwarding-policy {
    term one {
      from {
        protocol bgp;
        route-filter 10.1.0.0/16 orlonger;
      }
    }
  }
}
```



```

    }
    then {
        install-nexthop lsp mc-c-lsp-1;
        accept;
    }
}
term two {
    from {
        protocol bgp;
        route-filter 10.2.0.0/16 orlonger;
    }
    then {
        install-nexthop lsp mc-c-lsp-2;
        accept;
    }
}
term three {
    from {
        protocol bgp;
        route-filter 10.3.0.0/16 orlonger;
    }
    then {
        install-nexthop lsp mc-c-lsp-3;
        accept;
    }
}
}
}
protocols {
    mpls {
        label-switched-path mc-c-lsp-1 {
            from 10.10.20.101;
            to 10.10.20.103;
        }
        label-switched-path mc-c-lsp-2 {
            from 10.10.20.101;
            to 10.10.20.103;
        }
        label-switched-path mc-c-lsp-3 {
            from 10.10.20.101;
            to 10.10.20.103;
        }
    }
}
}

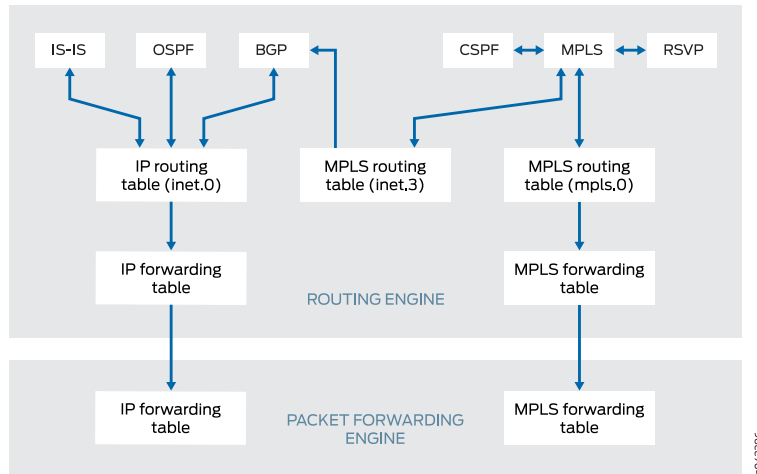
```

MPLS and Routing Tables

The IGP and BGP store their routing information in the `inet.0` routing table, the main IP routing table. If the **traffic-engineering bgp** command is configured, thereby allowing only BGP to use MPLS paths for forwarding traffic, MPLS path information is stored in a separate routing table, `inet.3`. Only BGP accesses the `inet.3` routing table. BGP uses both `inet.0` and `inet.3` to resolve next-hop addresses. If the **traffic-engineering bgp-igp** command is configured, thereby allowing the IGPs to use MPLS paths for forwarding traffic, MPLS path information is stored in the `inet.0` routing table. (Figure 14 on page 44 and

Figure 15 on page 45 illustrate the routing tables in the two traffic engineering configurations.)

Figure 14: Routing and Forwarding Tables, traffic-engineering bgp



The inet.3 routing table contains the host address of each LSP's egress router. This routing table is used on ingress routers to route packets to the destination egress router. BGP uses the inet.3 routing table on the ingress router to help in resolving next-hop addresses.

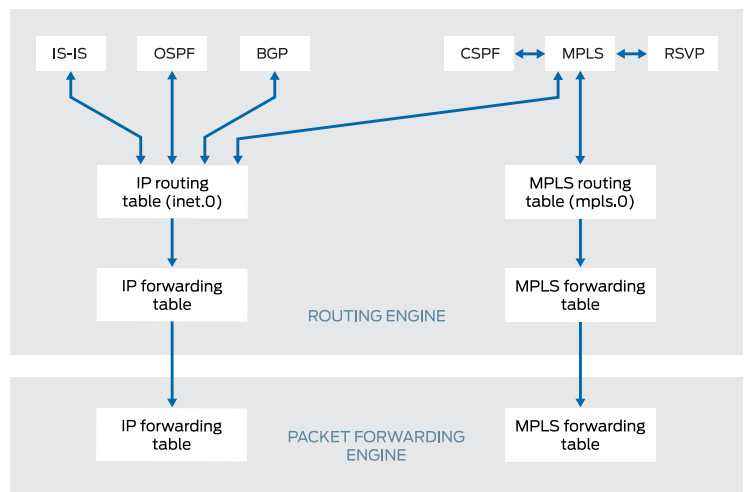
MPLS also maintains an MPLS path routing table (mpls.0), which contains a list of the next label-switched router in each LSP. This routing table is used on transit routers to route packets to the next router along an LSP.

Typically, the egress router in an LSP does not consult the mpls.0 routing table. (This router does not need to consult mpls.0 because the penultimate router in the LSP either changes the packet's label to a value of 0 or pops the label.) In either case, the egress router forwards it as an IPv4 packet, consulting the IP routing table, inet.0, to determine how to forward the packet.

When a transit or egress router receives an MPLS packet, information in the MPLS forwarding table is used to determine the next transit router in the LSP or to determine that this router is the egress router.

When BGP resolves a next-hop prefix, it examines both the inet.0 and inet.3 routing tables, seeking the next hop with the lowest preference. If it finds a next-hop entry with an equal preference in both routing tables, BGP prefers the entry in the inet.3 routing table.

Figure 15: Routing and Forwarding Tables, traffic-engineering bgp-igp



Generally, BGP selects next-hop entries in the inet.3 routing table because their preferences are always lower than OSPF and IS-IS next-hop preferences. When you configure LSPs, you can override the default preference for MPLS LSPs, which might alter the next-hop selection process.

When BGP selects a next-hop entry from the inet.3 routing table, it installs that LSP into the forwarding table in the Packet Forwarding Engine, which causes packets destined for that next hop to enter and travel along the LSP. If the LSP is removed or fails, the path is removed from the inet.3 routing table and from the forwarding table, and BGP reverts to using a next hop from the inet.0 routing table.

MPLS and Traffic Protection

Typically, when an LSP fails, the router immediately upstream from the failure signals the outage to the ingress router. The ingress router calculates a new path to the egress router, establishes the new LSP, and then directs the traffic from the failed path to the new path. This rerouting process can be time-consuming and prone to failure. For example, the outage signals to the ingress router might get lost, or the new path might take too long to come up, resulting in significant packet drops. The Junos OS provides several complementary mechanisms for protecting against LSP failures:

- Standby secondary paths—You can configure primary and secondary paths. You configure secondary paths with the **standby** statement. To activate traffic protection, you need to configure these standby paths only on the ingress router. If the primary path fails, the ingress router immediately reroutes traffic from the failed path to the standby path, thereby eliminating the need to calculate a new route and signal a new path. For information about configuring standby LSPs, see [“Configuring Hot Standby of Secondary Paths for LSPs” on page 267](#).
- Fast reroute—You configure fast reroute on an LSP to minimize the effect of a failure in the LSP. Fast reroute enables a router upstream from the failure to route around the failure quickly to the router downstream of the failure. The upstream router then signals

the outage to the ingress router, thereby maintaining connectivity before a new LSP is established. For a detailed overview of fast reroute, see [“Fast Reroute Overview” on page 46](#). For information about configuring fast reroute, see [“Configuring Fast Reroute” on page 226](#).

- **Link protection**—You can configure link protection to help ensure that traffic traversing a specific interface from one router to another can continue to reach its destination in the event that this interface fails. When link protection is configured for an interface and configured for an LSP that traverses this interface, a bypass LSP is created that handles this traffic if the interface fails. The bypass LSP uses a different interface and path to reach the same destination. For information about configuring link protection, see [“Configuring Link Protection on Interfaces Used by LSPs” on page 502](#).

When standby secondary path, and fast reroute or link protection are configured on an LSP, full traffic protection is enabled. When a failure occurs in an LSP, the router upstream from the failure routes traffic around the failure and notifies the ingress router of the failure. This rerouting keeps the traffic flowing while waiting for the notification to be processed at the ingress router. After receiving the failure notification, the ingress router immediately reroutes the traffic from the patched primary path to the more optimal standby path.

Fast reroute and link protection provide a similar type of traffic protection. Both features provide a quick transfer service and employ a similar design. Fast reroute and link protection are both described in RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*. However, you need to configure only one or the other. Although you can configure both, there is little, if any, benefit in doing so.


**Related
Documentation**

- [Configuring Hot Standby of Secondary Paths for LSPs on page 267](#)
- [Fast Reroute Overview on page 46](#)
- [Configuring Fast Reroute on page 226](#)
- [Configuring Link Protection on Interfaces Used by LSPs on page 502](#)

Fast Reroute Overview


Fast reroute provides redundancy for an LSP path. When you enable fast reroute, detours are precomputed and preestablished along the LSP. In case of a network failure on the current LSP path, traffic is quickly routed to one of the detours. [Figure 16 on page 47](#) illustrates an LSP from Router A to Router F, showing the established detours. Each detour is established by an upstream node to avoid the link toward the immediate downstream node and the immediate downstream node itself. Each detour might traverse through one or more label-switched routers (or switches) that are not shown in the figure.

Fast reroute protects traffic against any single point of failure between the ingress and egress routers (or switches). If there are multiple failures along an LSP, fast reroute itself might fail. Also, fast reroute does not protect against failure of the ingress or egress routers.

Figure 16: Detours Established for an LSP Using Fast Reroute

If a node detects that a downstream link has failed (using a link-layer-specific liveness detection mechanism) or that a downstream node has failed (for example, using the RSVP neighbor hello protocol), the node quickly switches the traffic to the detour and, at the same time, signals the ingress router about the link or node failure.

Figure 17 on page 47 illustrates the detour taken when the link between Router B and Router C fails.

Figure 17: Detour After the Link from Router B to Router C Fails

If the network topology is not rich enough (there are not enough routers with sufficient links to other routers), some of the detours might not succeed. For example, the detour from Router A to Router C in Figure 16 on page 47 cannot traverse link A-B and Router B. If such a path is not possible, the detour does not occur.

Note that after the node switches traffic to the detour, it might switch the traffic again to a newly calculated detour soon after. This is because the initial detour route might not be the best route. To make rerouting as fast as possible, the node switches traffic onto the initial detour without first verifying that the detour is valid. Once the switch is made, the node recomputes the detour. If the node determines that the initial detour is still valid, traffic continues to flow over this detour. If the node determines that the initial detour is no longer valid, it again switches the traffic to a newly computed detour.



NOTE: If you issue `show` commands after the node has switched traffic to the initial detour, the node might indicate that the traffic is still flowing over the original LSP. This situation is temporary and should correct itself quickly.

The time required for a fast-rerouting detour to take effect depends on two independent time intervals:

- Amount of time to detect that there is a link or node failure—This interval depends greatly on the link layer in use and the nature of the failure. For example, failure detection

on an SONET/SDH link typically is much faster than on a Gigabit Ethernet link, and both are much faster than detection of a router failure.


- Amount of time required to splice the traffic onto the detour—This operation is performed by the Packet Forwarding Engine, which requires little time to splice traffic onto the detour. The time needed can vary depending on the number of LSPs being switched to detours.

Fast reroute is a short-term patch to reduce packet loss. Because detour computation might not reserve adequate bandwidth, the detours might introduce congestion on the alternate links. The ingress router is the only router that is fully aware of LSP policy constraints and, therefore, is the only router able to come up with adequate long-term alternate paths.

Detours are created by use of RSVP and, like all RSVP sessions, they require extra state and overhead in the network. For this reason, each node establishes at most one detour for each LSP that has fast reroute enabled. Creating more than one detour for each LSP increases the overhead, but serves no practical purpose.

To reduce network overhead further, each detour attempts to merge back into the LSP as soon as possible after the failed node or link. If you can consider an LSP that travels through n router nodes, it is possible to create $n - 1$ detours. For instance, in [Figure 18 on page 48](#), the detour tries to merge back into the LSP at Router D instead of at Router E or Router F. Merging back into the LSP makes the detour scalability problem more manageable. If topology limitations prevent the detour from quickly merging back into the LSP, detours merge with other detours automatically.

Figure 18: Detours Merging into Other Detours



Related Documentation

- [fast-reroute on page 870](#)
- [Configuring Fast Reroute on page 226](#)
- *MPLS Feature Support on QFX Series and EX4600 Switches*
- *Interprovider and Carrier-of-Carriers VPNs*

Detour Merging Process

This section describes the process used by a router to determine which LSP to select when the router receives path messages from different interfaces with identical Session and Sender Template objects. When this occurs, the router needs to merge the path states.

The router employs the following process to determine when and how to merge path states:

- When all the path messages do not include a fast reroute or a detour object, or when the router is the egress of the LSP, no merging is required. The messages are processed according to RSVP traffic engineering.
- Otherwise, the router *must* record the path state in addition to the incoming interface. If the path messages do not share the same outgoing interface and next-hop router, the router considers them to be independent LSPs and does not merge them.
- For all the path messages that share the same outgoing interface and next-hop router, the router uses the following process to select the final LSP:
 - If only one LSP originates from this node, select it as the final LSP.
 - If only one LSP contains a fast reroute object, select it as the final LSP.
 - If there are several LSPs and some of them have a detour object, eliminate those containing a detour object from the final LSP selection process.
 - If several final LSP candidates remain (that is, there are still both detour and protected LSPs), select the LSPs with fast reroute objects.
 - If none of the LSPs have fast reroute objects, select the ones without detour objects. If all the LSPs have detour objects, select them all.
 - Of the remaining LSP candidates, eliminate from consideration those that traverse nodes that other LSPs avoid.
 - If several candidate LSPs still remain, select the one with the shortest explicit route object (ERO) path length. If more than one LSP has the same path length, select one randomly.
- Once the final LSP has been identified, the router must transmit only the path messages that correspond to this LSP. All other LSPs are considered merged at this node.

Detour Computations

Computing and setting up detours is done independently at each node. On a node, if an LSP has fast reroute enabled and if a downstream link or node can be identified, the router performs a Constrained Shortest Path First (CSPF) computation using the information in the local traffic engineering database. For this reason, detours rely on your IGP supporting traffic engineering extensions. Without the traffic engineering database, detours cannot be established.

CSPF initially attempts to find a path that skips the next downstream node. Attempting to find this path provides protection against downstream failures in either nodes or links. If a node-skipping path is not available, CSPF attempts to find a path on an alternate link to the next downstream node. Attempting to find an alternate link provides protection against downstream failures in links only. Detour computations might not succeed the first time. If a computation fails, the router recomputes detours approximately once every refresh interval until the computation succeeds. The RSVP metric for each detour is set to a value in the range from 10,000 through 19,999.

Fast Reroute Path Optimization

A fast reroute protection path is nondeterministic. The actual protection path of a particular node depends on the history of the LSP and the network topology when the fast reroute path was computed. The lack of deterministic behavior can lead to operational difficulties and poorly optimized paths after multiple link flaps in a network. Even in a small network, after a few link flaps fast reroute paths can traverse an arbitrarily large number of nodes and can remain in that state indefinitely. This is inefficient and makes the network less predictable.

Fast reroute optimization addresses this deficiency. It provides a global path optimization timer, allowing you to optimize all LSPs that have fast reroute enabled and a detour path up and running. The timer value can be varied depending on the expected RE processing load.

The fast reroute optimization algorithm is based on the IGP metric only. As long as the new path's IGP metric is lower than the old path's, the CSPF result is accepted, even if the new path might be more congested (higher bandwidth utilization) or traverses more hops.

In conformance with RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*, when a new path is computed and accepted for fast reroute optimization, the existing detour is destroyed first and then the new detour is established. To prevent traffic loss, detours actively protecting traffic are not optimized.

On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview

This topic describes methods for measuring packet loss, delay, and throughput for point-to-point ultimate hop popping (UHP) label-switched paths (LSPs) in MPLS networks to enable monitoring of network performance.

- [Importance of Measuring Packet Loss and Delay on page 50](#)
- [Defining Packet Loss, Delay, and Throughput on page 51](#)
- [Packet Loss and Delay Measurement Mechanisms on page 51](#)
- [Packet Loss and Delay Metrics on page 52](#)
- [Packet Loss and Delay Measurement Concepts on page 52](#)
- [Packet Loss and Delay Measurement Functionality on page 55](#)
- [Packet Loss and Delay Features on page 56](#)

Importance of Measuring Packet Loss and Delay

The rise of bandwidth-consuming applications, such as IPTV and mobile video, coupled with the pressure to minimize the cost per bit and maximize the value per bit, is forcing carriers to transition their transport networks from circuit-based technologies to packet-based technologies. MPLS is a widely successful, connection-oriented packet transport technology that is ideally suited for packet-based transport networks.

With the emergence of new applications on data networks, it is becoming increasingly important for service providers to accurately predict the impact of new application rollouts. Understanding and modelling network performance in the network is especially relevant for deployment of new-world applications to ensure successful implementations. In packet networks, packet loss and delay are two of the most fundamental measures of performance. Their role is even more central when it comes to end-to-end measurements.

The traffic belonging to most of the end-to-end user applications is either loss sensitive (file transfer), delay sensitive (voice or video applications), or both (interactive computing applications). The service-level agreements (SLAs) of service providers depend on the ability to measure and monitor these network performance metrics, as the SLAs are directly or indirectly dependent on the loss and delay the customer traffic experiences in the service provider network.

To ensure compliance to the SLA, service providers need tools to measure and monitor the performance metrics for packet loss, one-way delay and two-way delay, and related metrics, such as delay variation and channel throughput. This measurement capability provides service providers with greater visibility into the performance characteristics of their networks, thereby facilitating planning, troubleshooting, and network performance evaluation.

Defining Packet Loss, Delay, and Throughput

In packet networks, packet loss and delay are two of the most fundamental measures of performance.

- **Loss**—Packet loss is the failure of one or more transmitted packets to arrive at their destination. Packet loss refers to the packets of data that are dropped by the network to manage congestion.

Data applications are very tolerant to packet loss, as they are generally not time sensitive and can retransmit the packets that were dropped. However, in video conference environments and pure audio communications, such as VoIP, packet loss can create jitter.

- **Delay**—Packet delay (also called latency) is the amount of time it takes for a packet of data to get from one designated point to another, depending on the speed of the transmission medium, such as copper wire, optical fiber, or radio waves, and the delays in transmission by devices along the way, such as routers and modems.

A low latency indicates a high network efficiency.

- **Throughput**—Packet delay measures the amount of time between the start of an action and its completion, whereas throughput is the total number of such actions that occur in a given amount of time.

Packet Loss and Delay Measurement Mechanisms

Packet delay and loss are two fundamental measures of network performance. Junos OS provides an on-demand mechanism to measure packet loss and delay over associated bidirectional MPLS ultimate hop popping (UHP) label-switched paths (LSPs).

The on-demand delay and packet loss measurement mechanism is initiated using the following CLI commands:

- **monitor mpls loss rsvp**—Performs an on-demand loss measurement for associated bidirectional UHP LSPs.
- **monitor mpls delay rsvp**—Performs an on-demand delay measurement for associated bidirectional UHP LSPs.
- **monitor mpls loss-delay rsvp**—Performs an on-demand combined loss and delay measurement for associated bidirectional UHP LSPs.

For initiating the delay and packet loss measuring mechanism, the desired parameters for measurement, such as the type of measurement and LSP name, need to be entered. On receiving the parameters, a summary of the performance monitoring data is displayed and the mechanism is terminated.

Packet Loss and Delay Metrics

The following performance metrics are measured using the on-demand packet loss and delay mechanisms:

- Loss measurement (packet and octet)
- Throughput measurement (packet and octet)
- Two-way channel delay
- Round-trip delay
- Inter-packet delay variation (IPDV)

The **monitor mpls loss rsvp** command performs the loss and throughput measurement, and the **monitor mpls delay rsvp** command performs the two-way channel delay, round-trip delay, and IPDV measurements. The **monitor mpls loss-delay rsvp** command performs a combined loss and delay measurement and measures all of the above-mentioned performance metrics simultaneously.

Packet Loss and Delay Measurement Concepts

The following concepts help to better understand the functionality of packet loss and delay:

- **Querier**—A querier is the ingress provider edge (PE) router, which originates the query message for loss or delay measurement.
- **Responder**—A responder is the egress PE router, which receives and responds to the query messages from a querier.
- **Associated bidirectional LSP**—An associated bidirectional LSP consists of two unidirectional LSPs that are tied together (or associated with each other) through configuration on both of the LSP end points.

The on-demand loss and delay measurement can be carried out only on associated bidirectional UHP LSPs.

- **Generic associated channel (G-Ach)**—The performance monitoring messages for the on-demand loss and delay measurement flow over the MPLS G-Ach. This type of channel supports only in-band responses, and does not provide support for out-of-band or no-response modes.
- **Measurement point (MP)**—MP is the location at which a condition is described for the measurement.

The MP for packet loss on the transmit side is between the switching fabric and the transmit interface. The counter value is stamped in the loss measurement message in the hardware before it is queued for transmission.

The MP for packet loss on the receive side is between the receive interface and the switching fabric. The MP is distributed on the receive side. Furthermore, when the transmit interface is an aggregate interface, the MP is distributed as well.

- **Query rate**—Query rate is the interval between two queries sent for loss and delay measurement.

Because the loss and delay measurement messages originate from the Routing Engine, a high query rate for multiple channels puts a heavy burden on the Routing Engine. The minimum query interval supported is 1 second.

The query rate should be high for 32-bit counters, because the counters might wrap quickly when data traffic rate is very high. The query rate can be low when 64-bit counters are in use at all the four measurement point locations involved in loss measurement. Junos OS supports only 64-bit counters.

- **Traffic class**—By default, loss measurement is supported for the whole channel. Junos OS also supports traffic class scoped packet loss measurement, where counters that maintain data traffic statistics per traffic class have to be created.

Per traffic class counters are not created by default. To configure traffic class scoped loss measurement, include the **traffic-class-statistics** statement at the **[edit protocols mpls statistics]** hierarchy level.

When **traffic-class-statistics** is configured, control packets flowing over the G-Ach are not counted in the transmit and receive counters.



NOTE: Enabling and disabling of traffic class statistics results in the resetting of all counters (aggregate counter and per-class counters) for the LSPs.

- **Loss measurement mode**—Junos OS supports the direct-mode of on-demand loss measurement, and does not provide support for the inferred-mode.

Direct loss measurement requires data traffic statistics to be maintained at the ingress and egress of two unidirectional LSPs of the associated bidirectional LSP. When an MX Series router is using only MPCs and MICs, counters to maintain data traffic statistics are created by default at the ingress of all types of LSPs and egress of UHP LSPs.

However, the direct-mode of loss measurement is not fully accurate due to the following reasons:

- Parallel forwarding nature of the hardware.
- Presence of equal cost multipath (ECMP) in the network, such as aggregated Ethernet interfaces, which can result in re-ordering of data packets relative to the loss measurement messages.
- Control packets that do not flow over G-Ach are not counted at the LSP ingress, but are counted at the LSP egress.
- Data traffic re-ordering relative to the loss measurement message when a Diffserv is implemented in the MPLS network and loss measurement scope is the complete channel and not traffic class scoped.

To overcome this limitation, perform traffic class scoped loss measurement when a Diffserv is implemented.



NOTE: Direct mode loss measurement is vulnerable to disruption when the ingress or egress interface associated with the LSP changes.

- **Loss measurement synchronization**—The synchronization conditions specified in section 2.9.8 of RFC 6374 do not hold true in the absolute sense. However, as the loss measurement counters are stamped in hardware, the errors introduced due to not satisfying the synchronization conditions are relatively small. These errors need to be quantified.

When the transmit or receive interface of the LSP is an aggregate interface, more errors are introduced as compared to when the interfaces are non-aggregate interfaces. In any case, the loss measurement counters are stamped in hardware, and the error needs to be quantified.

- **Delay measurement accuracy**—When the transmit and receive interfaces reside on different Packet Forwarding Engines, the clock must be synchronized on these Packet Forwarding Engines for two-way delay measurements. This condition holds true for the platform on which the on-demand delay measurement feature is implemented.

When there are aggregate interfaces or ECMP, the delay is measured for only one of the potential paths.

When a combined loss and delay message is used for delay calculation, the accuracy of delay is lower compared to when the delay measurement message is used in some cases, such as when the transmit or receive interface is an aggregate interface.


Delay measurement is always performed on a per-traffic-class basis, and the accuracy of the measurement needs to be quantified after testing.

- **Timestamp format**—Junos OS supports only the IEEE 1588 Precision Time Protocol (PTP) [IEEE1588] format for recording delay measurement messages. Network Time Format (NTP) is not supported.
- **Operations, administration, and maintenance (OAM)**—To indicate that all the OAM messages for MPLS LSPs flow over the MPLS G-Ach, and to enable the MPLS performance monitoring messages to be carried over the MPLS G-Ach, the **oam mpls-tp-mode** statement must be included at the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level.

Packet Loss and Delay Measurement Functionality

Figure 19 on page 55 illustrates the basic methods used for the bidirectional measurement of packet loss and delay. A bidirectional channel exists between the two routers, Router A and Router B. The temporal reference points – T1, T2, T3, and T4 – are associated with a measurement operation that takes place at Router A. The operation consists of Router A sending a query message to Router B, and Router B sending back a response. Each reference point indicates the point of time at which either the query or the response message is transmitted or received over the channel.

Figure 19: Basic Bidirectional Measurement



In Figure 19 on page 55, Router A can arrange to measure the packet loss over the channel in the forward and reverse directions by sending loss measurement query messages to Router B. Each of the forward and reverse messages contain the count of packets transmitted prior to time T1 over the channel to Router B (A_TxP).

When the message reaches Router B, two values are appended to the message and the message is reflected back to Router A. The two values are the count of packets received prior to time T2 over the channel from Router A (B_RxP) and the count of packets transmitted prior to time T3 over the channel to Router A (B_TxP).

When the response reaches Router A, a fourth value is appended to the message – the count of packets received prior to time T4 over the channel from Router B (A_RxP).

These four counter values – (A_TxP), (B_RxP), (B_TxP), and (A_RxP) – enable Router A to compute the desired loss statistics. Because the transmit count at Router A and the receive count at Router B (and vice versa) might not be synchronized at the time of the first message, and to limit the effects of counter wrap, the loss is computed in the form of a delta between the messages.

The transmit loss (A_TxLoss[n-1,n]) and receive loss (A_RxLoss[n-1,n]) within the measurement interval marked by the messages LM[n-1] and LM[n] are computed by Router A as follows:

$$A_TxLoss[n-1,n] = (A_TxP[n] - A_TxP[n-1]) - (B_RxP[n] - B_RxP[n-1])$$

$$A_RxLoss[n-1,n] = (B_TxP[n] - B_TxP[n-1]) - (A_RxP[n] - A_RxP[n-1])$$

The arithmetic is modulo the counter size.

To measure at Router A the delay over the channel to Router B, a delay measurement query message is sent from Router A to Router B containing a timestamp recording the instant at which it is transmitted. In [Figure 19 on page 55](#), the timestamp is recorded in T1.

When the message reaches Router B, a timestamp is added, recording the instant at which it is received (T2). The message can now be reflected from Router B to Router A, with Router B adding its transmit timestamp (T3) and Router A adding its receive timestamp (T4).

These four timestamps – T1, T2, T3, and T4 – enable Router A to compute the one-way delay in each direction, as well as the two-way delay for the channel. The one-way delay computations require that the clocks of Routers A and B be synchronized.

At this point, Router A can compute the two-way channel delay and round-trip delay associated with the channel as follows:

$$\text{Two-way channel delay} = (T4 - T1) - (T3 - T2)$$

$$\text{Round-trip delay} = T4 - T1$$

Packet Loss and Delay Features

Supported Features of Packet Loss and Delay

Junos OS supports the following features with on-demand loss and delay measurement:

- Performance monitoring for associated bidirectional MPLS point-to-point UHP LSPs only
- Loss measurement
- Throughput measurement
- Two-way delay measurement (channel delay and round-trip delay)
- Inter-packet delay variation (IPDV)
- Direct-mode loss measurement
- Aggregated Ethernet and aggregated SONET interfaces
- Multichassis support
- 64-bit compatible

Unsupported Features of Packet Loss and Delay

Junos OS does not support the following on-demand loss and delay measurement functionality:

- Loss and delay measurement for pseudowires (section 2.9.1 of RFC 6374)
- Unidirectional measurement (section 2.6 of RFC 6374)
- Dyadic measurement (section 2.7 of RFC 6374)
- Loss and delay measurement in loopback mode (section 2.8 of RFC 6374)
- Loss and delay measurement to an intermediate node from an LSP endpoint (section 2.9.5 of RFC 6374)
- External post-processing (section 2.9.7 of RFC 6374)
- Inferred-mode loss measurement (section 2.9.8 of RFC 6374)
- Pro-active mode
- Logical systems
- SNMP

**Related
Documentation**

- [Example: Configuring On-Demand Loss and Delay Measurement on page 439](#)
- [monitor mpls loss rsvp on page 1154](#)
- [monitor mpls delay rsvp on page 1150](#)
- [monitor mpls loss-delay rsvp on page 1159](#)

CHAPTER 3

Configuring MPLS Routers

- [Minimum MPLS Configuration on page 59](#)
- [Configuring the Ingress Router for MPLS-Signaled LSPs on page 60](#)
- [Example: Configuring an Explicit-Path LSP on page 65](#)
- [Example: Configuring a Constrained-Path LSP for Which Junos OS Makes All Forwarding Decisions on page 66](#)
- [Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most Forwarding Decisions and Considers Hop Constraints on page 66](#)
- [Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most Forwarding Decisions and the Secondary Path Is Explicit on page 67](#)
- [Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs on page 68](#)
- [Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages on page 68](#)
- [Configuring MPLS-Signaled LSPs to Use GRE Tunnels on page 70](#)
- [Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks on page 71](#)
- [SRLG Overview on page 80](#)
- [Example: Configuring SRLG on page 81](#)
- [Example: Excluding SRLG Links Completely for the Secondary LSP on page 90](#)
- [Example: Configuring SRLG with Link Protection on page 95](#)
- [Example: Configuring SRLG with Link Protection with the exclude-srlg Option on page 116](#)
- [Configuring the MPLS Transport Profile for OAM on page 136](#)
- [Configuring MPLS Inter-AS Link-Node Protection with Labeled BGP on page 148](#)
- [Configuring Egress Protection Service Mirroring for BGP Signaled Layer 2 Services on page 164](#)
- [Example: Configuring MPLS Egress Protection Service Mirroring for BGP Signaled Layer 2 Services on page 168](#)
- [Example: Configuring Layer 3 VPN Egress Protection with PLR as Protector on page 184](#)

Minimum MPLS Configuration

To enable MPLS on the router, you must include at least the following statements. This minimum configuration enables MPLS on a logical interface. All other MPLS configuration

statements are optional. Note that this configuration does nothing more than enable MPLS on the router and on the specified interface. It could allow RSVP-signaled MPLS traffic to transit the router.

Include the **family mpls** statement:

```
family mpls;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Include the interface in the MPLS and RSVP protocol configuration:

```
mpls {  
  interface (interface-name | all); # Required to enable MPLS on the interface  
}  
rsvp { # Required for RSVP-signaled MPLS only  
  interface interface-name;  
}
```

You can configure these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

For every interface you enable, two special routes are installed automatically in the MPLS forwarding table. One route has a label value of 0, and the second has a label value of 1. (For information about these labels, see [“Special MPLS Labels” on page 25.](#))

Configuring the Ingress Router for MPLS-Signaled LSPs

MPLS-signaled label-switched paths (LSPs) run from a specific ingress router to a specific egress router. For basic MPLS-signaled LSP function, you must configure the ingress router, but do not have to configure any other routers.

To configure signaled LSPs, perform the following tasks on the ingress router:

- [Creating Named Paths on page 60](#)
- [Configuring Alternate Backup Paths Using Fate Sharing on page 62](#)

Creating Named Paths

To configure signaled LSPs, you must first create one or more named paths on the ingress router. For each path, you can specify some or all transit routers in the path, or you can leave it empty.

Each pathname can contain up to 32 characters and can include letters, digits, periods, and hyphens. The name must be unique within the ingress router. Once a named path is created, you can use the named path with the **primary** or **secondary** statement to configure

LSPs at the `[edit protocols mpls label-switched-path label-path-name]` hierarchy level. You can specify the same named path on any number of LSPs.

To determine whether an LSP is associated with the primary or secondary path in an RSVP session, issue the `show rsvp session detail` command.

To create an empty path, create a named path by including the following form of the `path` statement. This form of the `path` statement is empty, which means that any path between the ingress and egress routers is accepted. In actuality, the path used tends to be the same path as is followed by destination-based, best-effort traffic.

```
path path-name;
```

You can include this statement at the following hierarchy levels:

- `[edit protocols mpls]`
- `[edit logical-systems logical-system-name protocols mpls]`

To create a path in which you specify some or all transit routers in the path, include the following form of the `path` statement, specifying one address for each transit router:

```
path path-name {
  (address | hostname) <strict | loose>;
}
```

You can include this statement at the following hierarchy levels:

- `[edit protocols mpls]`
- `[edit logical-systems logical-system-name protocols mpls]`

In this form of the `path` statement, you specify one or more transit router addresses. Specifying the ingress or egress routers is optional. You can specify the address or hostname of each transit router, although you do not need to list each transit router if its type is `loose`. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path up to the egress router (optional) or the router immediately before the egress router. You need to specify only one address per router hop. If you specify more than one address for the same router, only the first address is used; the additional addresses are ignored and truncated.

For each router address, you specify the type, which can be one of the following:

- **strict**—(Default) The route taken from the previous router to this router is a direct path and cannot include any other routers. If **address** is an interface address, this router also ensures that the incoming interface is the one specified. Ensuring that the incoming interface is the one specified is important when there are parallel links between the previous router and this router. It also ensures that routing can be enforced on a per-link basis.

For strict addresses, you must ensure that the router immediately preceding the router you are configuring has a direct connection to that router. The address can be a loopback interface address, in which case the incoming interface is not checked.

- **loose**—The route taken from the previous router to this router need not be a direct path, can include other routers, and can be received on any interface. The address can be any interface address or the address of the loopback interface.

Examples: Creating Named Paths

Configure a path, **to-hastings**, to specify the complete strict path from the ingress to the egress routers through 14.1.1.1, 13.1.1.1, 12.1.1.1, and 11.1.1.1, in that order. There cannot be any intermediate routers except the ones specified. However, there can be intermediate routers between 11.1.1.1 and the egress router because the egress router is not specifically listed in the **path** statement. To prevent intermediate routers before egress, configure the egress router as the last router, with a **strict** type.

```
[edit protocols mpls]
path to-hastings {
  14.1.1.1 strict;
  13.1.1.1 strict;
  12.1.1.1 strict;
  11.1.1.1 strict;
}
```

Create a path, **alt-hastings**, to allow any number of intermediate routers between routers 14.1.1.1 and 11.1.1.1. In addition, intermediate routers are permitted between 11.1.1.1 and the egress router.

```
[edit protocols mpls]
path alt-hastings {
  14.1.1.1 strict;
  11.1.1.1 loose;
}
```

Configuring Alternate Backup Paths Using Fate Sharing

You can create a database of information that Constrained Shortest Path First (CSPF) uses to compute one or more backup paths in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. Because these network elements share the same fate, this relationship is called fate sharing.

You can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible to ensure that, if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination.

For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.

The following sections describe how to configure fate sharing and how it affects CSPF, and provides a fate sharing configuration example:

- [Configuring Fate Sharing on page 63](#)
- [Implications for CSPF on page 64](#)

- [Implications for CSPF When Fate Sharing with Bypass LSPs on page 64](#)
- [Example: Configuring Fate Sharing on page 64](#)

Configuring Fate Sharing

To configure fate sharing, include the **fate-sharing** statement:

```
fate-sharing {
  group group-name {
    cost value;
    from address <to address>;
  }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Each fate-sharing group must have a name, which can be up to 32 characters long and can contain letters, digits, periods (.) and hyphens (-). You can define up to 512 groups.

Fate-sharing groups contain three types of objects:

- Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; **from 1.2.3.4 to 1.2.3.5** and **from 1.2.3.5 to 1.2.3.4** have the same meaning.
- Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface **192.168.200.0/24** has four routers attached to it, each router link is individually identified:

```
from 192.168.200.1; # LAN interface of router 1
from 192.168.200.2; # LAN interface of router 2
from 192.168.200.3; # LAN interface of router 3
from 192.168.200.4; # LAN interface of router 4
```

You can list the addresses in any order.

- A router node—Identified by its configured router ID.

All objects in a group share certain similarities. For example, you can define a group for all fibers that share the same fiber conduit, all optical channels that share the same fiber, all links that connect to the same LAN switch, all equipment that shares the same power source, and so on. All objects are treated as /32 host addresses.

For a group to be meaningful, it should contain at least two objects. You can configure groups with zero or one object; these groups are ignored during processing.

An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database

does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.

Implications for CSPF

When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path.

When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs:

1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search.
2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together.
3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive.
4. The router performs regular CSPF computation with the adjusted topology.

Implications for CSPF When Fate Sharing with Bypass LSPs

When fate sharing is enabled with link protection or link-node protection, CSPF operates as follows when calculating the bypass LSP path:

- CSPF identifies the fate-sharing groups that are associated with the primary LSP path. CSPF does this by identifying the immediate downstream link and immediate downstream nodes that the bypass is trying to protect. CSPF compiles group lists that contain the immediate downstream link and immediate downstream nodes.
- CSPF checks each link (from ingress to the immediate downstream node) in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group.
- CSPF identifies the downstream link that is not in the fate-shared path.

This calculation prevents bypasses from using the same physical link as the primary LSP path when viable alternatives are available.

Example: Configuring Fate Sharing

Configure fate-sharing groups **east** and **west**. Because **west** has no objects, it is ignored during processing.

```
[edit routing-options]
fate-sharing {
  group east {
    cost 20; # Optional, default value is 1
```

```

    from 1.2.3.4 to 1.2.3.5; # A point-to-point link
    from 192.168.200.1; # LAN interface
    from 192.168.200.2; # LAN interface
    from 192.168.200.3; # LAN interface
    from 192.168.200.4; # LAN interface
    from 10.168.1.220; # Router ID of a router node
    from 10.168.1.221; # Router ID of a router node
  }
  group west {
    .....
  }
}

```

Example: Configuring an Explicit-Path LSP

On the ingress router, create an explicit-path LSP, and specify the transit routers between the ingress and egress routers. In this configuration, no constrained-path computation is performed. For the primary path, all intermediate hops are strictly specified so that its route cannot change. The secondary path must travel through router 14.1.1.1 first, then take whatever route is available to reach the destination. The remaining route taken by the secondary path is typically the shortest path computed by the IGP.

```

[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
protocols {
  rsvp {
    interface so-0/0/0;
  }
  mpls {
    path to-hastings {
      14.1.1.1 strict;
      13.1.1.1 strict;
      12.1.1.1 strict;
      11.1.1.1 strict;
    }
    path alt-hastings {
      14.1.1.1 strict;
      11.1.1.1 loose; # Any IGP route is acceptable
    }
    label-switched-path hastings {
      to 11.1.1.1;
      hop-limit 32;
      bandwidth 10m; # Reserve 10 Mbps
      no-cspf; # do not perform constrained-path computation
      primary to-hastings;
      secondary alt-hastings;
    }
  }
  interface so-0/0/0;
}

```

```
}
```

Example: Configuring a Constrained-Path LSP for Which Junos OS Makes All Forwarding Decisions

On the ingress router, create a constrained-path LSP in which the Junos OS makes all the forwarding decisions. When the LSP is successfully set up, a route toward 10.1.1.1/32 is installed in the inet.3 table so that all BGP routes with matching BGP next-hop addresses can be forwarded through the LSP.

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
protocols {
  rsvp {
    interface so-0/0/0;
  }
  mpls {
    label-switched-path to-hastings {
      to 10.1.1.1;
    }
    interface so-0/0/0;
  }
}
```

Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most Forwarding Decisions and Considers Hop Constraints

On the ingress router, create a constrained-path LSP in which the Junos OS makes most of the forwarding decisions, taking into account the hop constraints listed in the **path** statements. The LSP is adaptive so that no bandwidth double-counting occurs on links shared by primary and secondary paths. To acquire the necessary link bandwidth, this LSP is allowed to preempt lower priority sessions. Finally, this path always keeps the secondary path in hot-standby state for quick failover.

```
[edit protocols]
mpls {
  path to-hastings {
    14.1.1.1 loose;
  }
  path alt-hastings {
    12.1.1.1 loose;
    11.1.1.1 strict;
  }
  label-switched-path hastings {
    to 11.1.1.1;
    bandwidth 10m; # Reserve 10 Mbps
    priority 0 0; # Preemptive, but not preemptable
    adaptive; # Set adaptivity
  }
}
```



```

primary to-hastings;
secondary alt-hastings {
    standby;
    bandwidth 1m; # Reserve only 1 Mbps for the secondary path
}
}
interface all;
}

```

Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most Forwarding Decisions and the Secondary Path Is Explicit

On the ingress router, create a constrained-path LSP in which the Junos OS makes most of the forwarding decisions for the primary path, subject to constraints of the path **to-hastings**, and in which the secondary path is an explicit path. The primary path must transit green or yellow links and must stay away from red links. The primary path is periodically recomputed and reoptimized. Finally, this path always keeps the secondary path in hot-standby state for quick failover.

When the LSP is up—either because the primary or secondary path is up, or because both paths are up—the prefix 16.0.0.0/8 is installed in the inet.3 table so that all BGP routes whose BGP next hop falls within that range can use the LSP. Also, the prefix 17/8 is installed in the inet.0 table so that BGP can resolve only its next hop through that prefix. The route also can be reached with the **traceroute** or **ping** command. These two routes are in addition to the 11.1.1.1/32 route.

```

[edit protocols]
mpls {
  admin-groups {
    green 1;
    yellow 2;
    red 3;
  }
  path to-hastings {
    14.1.1.1 loose;
  }
  path alt-hastings {
    14.1.1.1 strict;
    13.1.1.1 strict;
    12.1.1.1 strict;
    11.1.1.1 strict;
  }
  label-switched-path hastings {
    to 11.1.1.1;
    bandwidth 100m;
    install 16.0.0.0/8; # in inet.3; cannot use to traceroute or ping
    install 17.0.0.0/8 active; # installed in inet.0; can use to traceroute or ping
    primary to-hastings {
      admin-group { # further constraints for path computation
        include-all [ green yellow ];
        exclude red;
      }
    }
    optimize-timer 3600; # reoptimize every hour
  }
}

```

```
secondary alt-hastings {  
    standby;  
    no-cspf; # do not perform constrained-path computation  
}  
}  
interface all;
```

Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs

To configure signaled LSPs on all MPLS routers that should participate in MPLS, you need to enable MPLS and RSVP on these routers, as described in [“Minimum MPLS Configuration” on page 59](#) and [“Minimum RSVP Configuration” on page 471](#).

Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages

An essential element of RSVP-based traffic engineering is the traffic engineering database. The traffic engineering database contains a complete list of all network nodes and links participating in traffic engineering, and a set of attributes each of those links can hold. (For more information about the traffic engineering database, see [“Constrained-Path LSP Computation” on page 29](#).) One of the most important link attributes is bandwidth.

Bandwidth availability on links changes quickly as RSVP LSPs are established and terminated. It is likely that the traffic engineering database will develop inconsistencies relative to the real network. These inconsistencies cannot be fixed by increasing the rate of IGP updates.

Link availability can share the same inconsistency problem. A link that becomes unavailable can break all existing RSVP LSPs. However, its unavailability might not readily be known by the network.

When you configure the **rsvp-error-hold-time** statement, a source node (ingress of an RSVP LSP) learns from the failures of its LSP by monitoring PathErr messages transmitted from downstream nodes. Information from the PathErr messages is incorporated into subsequent LSP computations, which can improve the accuracy and speed of LSP setup. Some PathErr messages are also used to update traffic engineering database bandwidth information, reducing inconsistencies between the traffic engineering database and the network.

You can control the frequency of IGP updates by using the **update-threshold** statement. See [“Configuring the RSVP Update Threshold on an Interface” on page 476](#).

This section discusses the following topics:

- [PathErr Messages on page 68](#)
- [Identifying the Problem Link on page 69](#)
- [Configuring the Router to Improve Traffic Engineering Database Accuracy on page 69](#)

PathErr Messages

PathErr messages report a wide variety of problems by means of different code and subcode numbers. You can find a complete list of these PathErr messages in RFC 2205,

Resource Reservation Protocol (RSVP), Version 1, Functional Specification and RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*.

When you configure the **rsvp-error-hold-time** statement, two categories of PathErr messages, which specifically represent link failures, are examined:

- Link bandwidth is low for this LSP: Requested bandwidth unavailable—code 1, subcode 2

This type of PathErr message represents a global problem that affects all LSPs transiting the link. They indicate that the actual link bandwidth is lower than that required by the LSP, and that it is likely that the bandwidth information in the traffic engineering database is an overestimate.

When this type of error is received, the available link bandwidth is reduced in the local traffic engineering database, affecting all future LSP computations.

- Link unavailable for this LSP:
 - Admission Control failure—code 1, any subcode except 2
 - Policy Control failures—code 2
 - Service Preempted—code 12
 - Routing problem—no route available toward destination—code 24, subcode 5

These types of PathErr messages are generally pertinent to the specified LSP. The failure of this LSP does not necessarily imply that other LSPs could also fail. These errors can indicate maximum transfer unit (MTU) problems, service preemption (either manually initiated by the operator or by another LSP with a higher priority), that a next-hop link is down, that a next-hop neighbor is down, or service rejection because of policy considerations. It is best to route this particular LSP away from the link.

Identifying the Problem Link

Each PathErr message includes the sender's IP address. This information is propagated unchanged toward the ingress router. A lookup in the traffic engineering database can identify the node that originated the PathErr message.

Each PathErr message carries enough information to identify the RSVP session that triggered the message. If this is a transit router, it simply forwards the message. If this router is the ingress router (for this RSVP session), it has the complete list of all nodes and links the session should traverse. Coupled with the originating node information, the link can be uniquely identified.

Configuring the Router to Improve Traffic Engineering Database Accuracy

To improve the accuracy of the traffic engineering database, configure the **rsvp-error-hold-time** statement. When this statement is configured, a source node (ingress of an RSVP LSP) learns from the failures of its LSP by monitoring PathErr messages transmitted from downstream nodes. Information from the PathErr messages is incorporated into subsequent LSP computations, which can improve the accuracy and speed of LSP setup. Some PathErr messages also are used to update traffic engineering

database bandwidth information, reducing inconsistencies between the traffic engineering database and the network.

To configure how long MPLS should remember RSVP PathErr messages and consider them in CSPF computation, include the **rsvp-error-hold-time** statement:

```
rsvp-error-hold-time seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

The time can be a value from 1 to 240 seconds. The default is 25 seconds. Configuring a value of 0 disables the monitoring of PathErr messages.

Configuring MPLS-Signaled LSPs to Use GRE Tunnels

MPLS LSPs can use generic routing encapsulation (GRE) tunnels to cross routing areas, autonomous systems, and ISPs. Bridging MPLS LSPs over an intervening IP domain is possible without disrupting the outlying MPLS domain.

LSPs can reach any destination that the GRE tunnels can reach. MPLS applications can be deployed without requiring all transit nodes to support MPLS, or requiring all transit nodes to support the same label distribution protocols (LDP or RSVP). If you use CSPF, you must configure OSPF or IS-IS through the GRE tunnel. Traffic engineering is not supported over GRE tunnels; for example, you cannot reserve bandwidth or set priority or preemption.



.....

NOTE: Use the **no-control word** statement to disable the control word when the topology uses GRE as the connection mechanism between provider edge routers and one of the provider edge routers is an M Series Multiservice Edge Router.

.....

For more information about GRE tunnels, see the *Junos OS Services Interfaces Library for Routing Devices*.

Example: Configuring MPLS-Signaled LSPs to Use GRE Tunnels

To configure MPLS over GRE tunnels:

1. Enable **family mpls** under the GRE interface configuration:

```
[edit interfaces]
interface gr-1/2/0 {
  unit 0 {
    tunnel {
      source 192.168.1.1;
      destination 192.168.1.2;
    }
    family inet {
      address 5.1.1.1/30;
    }
  }
}
```

```

    }
    family iso;
    family mpls;
  }
}

```

2. Enable RSVP and MPLS over the GRE tunnel:

```

[edit protocols]
rsvp {
  interface gr-1/2/0.0;
}
mpls {
  ...
  interface gr-1/2/0.0;
}

```

3. Configure LSPs to travel through the GRE tunnel endpoint address:

```

[edit protocols]
mpls {
  label-switched-path gre-tunnel {
    to 5.1.1.2;
    ...
  }
}

```

Standard LSP configuration options apply. If the routing table specifies that a particular route will traverse a GRE tunnel, the RSVP packets will traverse the tunnel as well.

Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks

This example shows how to configure the Junos OS to tunnel IPv6 over an MPLS-based IPv4 network. External BGP (EBGP) is used between the customer edge (CE) and provider edge (PE) devices. The remote CE devices have different AS numbers for loop detection.

- [Requirements on page 71](#)
- [Overview on page 71](#)
- [Configuration on page 74](#)
- [Verification on page 79](#)

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

Detailed information about the Juniper Networks implementation of IPv6 over MPLS is described in the following Internet drafts:

- Internet draft draft-ietf-l3vpn-bgp-ipv6-07.txt, *BGP-MPLS IP VPN extension for IPv6 VPN* (expires January 2006)


- Internet draft draft-ooms-v6ops-bgp-tunnel-06.txt, *Connecting IPv6 Islands over IPv4 MPLS using IPv6 Provider Edge Routers* (expires July 2006)

These Internet drafts are available on the IETF website at <http://www.ietf.org/>.

This example shows you how to interconnect a two IPv6 networks over an IPv4-based network core, giving you the ability to provide IPv6 service without having to upgrade the routers in your core network. Multiprotocol Border Gateway Protocol (MP-BGP) is configured to exchange routes between the IPv6 networks, and data is tunneled between these IPv6 networks by means of IPv4-based MPLS.

In [Figure 20 on page 72](#), Routers PE1 and PE2 are dual-stack BGP routers, meaning they have both IPv4 and IPv6 stacks. The PE routers link the IPv6 networks through the customer edge (CE) routers to the IPv4 core network. The CE routers and the PE routers connect through a link layer that can carry IPv6 traffic. The PE routers use IPv6 on the CE router-facing interfaces and use IPv4 and MPLS on the core-facing interfaces. Note that one of the connected IPv6 networks could be the global IPv6 Internet.

Figure 20: IPv6 Networks Linked by MPLS IPv4 Tunnels



The two PE routers are linked through an MP-BGP session using IPv4 addresses. They use the session to exchange IPv6 routes with an IPv6 (value 2) address family indicator (AFI) and a subsequent AFI (SAFI) (value 4). Each PE router sets the next hop for the IPv6 routes advertised on this session to its own IPv4 address. Because MP-BGP requires the BGP next hop to correspond to the same address family as the network layer reachability information (NLRI), this IPv4 address needs to be embedded within an IPv6 format.

The PE routers can learn the IPv6 routes from the CE routers connected to them using routing protocols Routing Information Protocol next generation (RIPng) or MP-BGP, or through static configuration. Note that if BGP is used as the PE-router-to-CE-router protocol, the MP-BGP session between the PE router and CE router could occur over an

IPv4 or IPv6 Transmission Control Protocol (TCP) session. Also, the BGP routes exchanged on that session would have SAFI unicast. You must configure an export policy to pass routes between IBGP and EBGP, and between BGP and any other protocol.

The PE routers have MPLS LSPs routed to each others' IPv4 addresses. IPv4 provides signaling for the LSPs by means of either LDP or RSVP. These LSPs are used to resolve the next-hop addresses of the IPv6 routes learned from MP-BGP. The next hops use IPv4-mapped IPv6 addresses, while the LSPs use IPv4 addresses.

The PE routers always advertise IPv6 routes to each other using a label value of 2, the explicit null label for IPv6 as defined in RFC 3032, *MPLS Label Stack Encoding*. As a consequence, each of the forwarding next hops for the IPv6 routes learned from remote PE routers normally push two labels. The inner label is 2 (this label could be different if the advertising PE router is not a Juniper Networks routing platform), and the outer label is the LSP label. If the LSP is a single-hop LSP, then only Label 2 is pushed.

It is also possible for the PE routers to exchange plain IPv6 routes using SAFI unicast. However, there is one major advantage in exchanging labeled IPv6 routes. The penultimate-hop router for an MPLS LSP can pop the outer label and then send the packet with the inner label as an MPLS packet. Without the inner label, the penultimate-hop router would need to discover whether the packet is an IPv4 or IPv6 packet to set the protocol field in the Layer 2 header correctly.

When the PE1 router in [Figure 20 on page 72](#) receives an IPv6 packet from the CE1 router, it performs a lookup in the IPv6 forwarding table. If the destination matches a prefix learned from the CE2 router, then no labels need to be pushed and the packet is simply sent to the CE2 router. If the destination matches a prefix that was learned from the PE2 router, then the PE1 router pushes two labels onto the packet and sends it to the provider router. The inner label is 2 and the outer label is the LSP label for the PE2 router.

Each provider router in the service provider's network handles the packet as it would any MPLS packet, swapping labels as it passes from provider router to provider router. The penultimate-hop provider router for the LSP pops the outer label and sends the packet to the PE2 router. When the PE2 router receives the packet, it recognizes the IPv6 explicit null label on the packet (Label 2). It pops this label and treats it as an IPv6 packet, performing a lookup in the IPv6 forwarding table and forwarding the packet to the CE3 router.

This example includes the following settings:

- In addition to configuring the **family inet6** statement on all the CE router-facing interfaces, you must also configure the statement on all the core-facing interfaces running MPLS. Both configurations are necessary because the router must be able to process any IPv6 packets it receives on these interfaces. You should not see any regular IPv6 traffic arrive on these interfaces, but you will receive MPLS packets tagged with Label 2. Even though Label 2 MPLS packets are sent in IPv4, these packets are treated as native IPv6 packets.
- You enable IPv6 tunneling by including the **ipv6-tunneling** statement in the configuration for the PE routers. This statement allows IPv6 routes to be resolved over an MPLS network by converting all routes stored in the inet.3 routing table to IPv4-mapped IPv6

addresses and then copying them into the inet6.3 routing table. This routing table can be used to resolve next hops for both inet6 and inet6-vpn routes.



NOTE: BGP automatically runs its import policy even when copying routes from a primary routing table group to a secondary routing table group. If IPv4 labeled routes arrive from a BGP session (for example, when you have configured the labeled-unicast statement at the [edit protocols bgp family inet] hierarchy level on the PE router), the BGP neighbor's import policy also accepts IPv6 routes, since the neighbor's import policy is run while doing the copy operation to the inet6.3 routing table.

- When you configure MP-BGP to carry IPv6 traffic, the IPv4 MPLS label is removed at the destination PE router. The remaining IPv6 packet without a label can then be forwarded to the IPv6 network. To enable this, include the **explicit-null** statement in the BGP configuration.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Device PE1

```
set interfaces fe-1/2/0 unit 2 family inet6 address ::10.1.1.2/126
set interfaces fe-1/2/0 unit 2 family mpls
set interfaces fe-1/2/1 unit 5 family inet address 10.1.1.5/30
set interfaces fe-1/2/1 unit 5 family inet6
set interfaces fe-1/2/1 unit 5 family mpls
set interfaces lo0 unit 2 family inet address 1.1.1.2/32
set protocols mpls ipv6-tunneling
set protocols mpls interface fe-1/2/0.2
set protocols mpls interface fe-1/2/1.5
set protocols bgp group toCE1 type external
set protocols bgp group toCE1 local-address ::10.1.1.2
set protocols bgp group toCE1 family inet6 unicast
set protocols bgp group toCE1 export send-bgp6
set protocols bgp group toCE1 peer-as 1
set protocols bgp group toCE1 neighbor ::10.1.1.1
set protocols bgp group toPE2 type internal
set protocols bgp group toPE2 local-address 1.1.1.2
set protocols bgp group toPE2 family inet6 labeled-unicast explicit-null
set protocols bgp group toPE2 export next-hop-self
set protocols bgp group toPE2 export send-v6
set protocols bgp group toPE2 neighbor 1.1.1.4
set protocols ospf area 0.0.0.0 interface fe-1/2/1.5
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols ldp interface fe-1/2/1.5
set policy-options policy-statement next-hop-self then next-hop self
set policy-options policy-statement send-bgp6 from family inet6
set policy-options policy-statement send-bgp6 from protocol bgp
set policy-options policy-statement send-bgp6 then accept
set policy-options policy-statement send-v6 from family inet6
```



```

set policy-options policy-statement send-v6 from protocol bgp
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.2
set routing-options autonomous-system 2

```

Device PE2

```

set interfaces fe-1/2/0 unit 10 family inet address 10.1.1.10/30
set interfaces fe-1/2/0 unit 10 family inet6
set interfaces fe-1/2/0 unit 10 family mpls
set interfaces fe-1/2/1 unit 13 family inet6 address ::10.1.1.13/126
set interfaces fe-1/2/1 unit 13 family mpls
set interfaces lo0 unit 4 family inet address 1.1.1.4/32
set protocols mpls ipv6-tunneling
set protocols mpls interface fe-1/2/0.10
set protocols mpls interface fe-1/2/1.13
set protocols bgp group toPE1 type internal
set protocols bgp group toPE1 local-address 1.1.1.4
set protocols bgp group toPE1 family inet6 labeled-unicast explicit-null
set protocols bgp group toPE1 export next-hop-self
set protocols bgp group toPE1 export send-v6
set protocols bgp group toPE1 neighbor 1.1.1.2
set protocols bgp group toCE3 type external
set protocols bgp group toCE3 local-address ::10.1.1.13
set protocols bgp group toCE3 family inet6 unicast
set protocols bgp group toCE3 export send-bgp6
set protocols bgp group toCE3 peer-as 3
set protocols bgp group toCE3 neighbor ::10.1.1.14
set protocols ospf area 0.0.0.0 interface fe-1/2/0.10
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols ldp interface fe-1/2/0.10
set policy-options policy-statement next-hop-self then next-hop self
set policy-options policy-statement send-bgp6 from family inet6
set policy-options policy-statement send-bgp6 from protocol bgp
set policy-options policy-statement send-bgp6 then accept
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol bgp
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.4
set routing-options autonomous-system 2

```

Device P

```

set interfaces fe-1/2/0 unit 6 family inet address 10.1.1.6/30
set interfaces fe-1/2/0 unit 6 family inet6
set interfaces fe-1/2/0 unit 6 family mpls
set interfaces fe-1/2/1 unit 9 family inet address 10.1.1.9/30
set interfaces fe-1/2/1 unit 9 family inet6
set interfaces fe-1/2/1 unit 9 family mpls
set interfaces lo0 unit 3 family inet address 1.1.1.3/32
set protocols mpls interface fe-1/2/0.6
set protocols mpls interface fe-1/2/1.9
set protocols ospf area 0.0.0.0 interface fe-1/2/0.6
set protocols ospf area 0.0.0.0 interface fe-1/2/1.9
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols ldp interface fe-1/2/0.6
set protocols ldp interface fe-1/2/1.9

```

```
set routing-options router-id 1.1.1.3
set routing-options autonomous-system 2
```

Device CE1

```
set interfaces fe-1/2/0 unit 1 family inet6 address ::10.1.1.1/126
set interfaces fe-1/2/0 unit 1 family mpls
set interfaces lo0 unit 1 family inet6 address ::1.1.1.1/128
set protocols bgp group toPE1 type external
set protocols bgp group toPE1 local-address ::10.1.1.1
set protocols bgp group toPE1 family inet6 unicast
set protocols bgp group toPE1 export send-v6
set protocols bgp group toPE1 peer-as 2
set protocols bgp group toPE1 neighbor ::10.1.1.2
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.1
set routing-options autonomous-system 1
```

Device CE3

```
set interfaces fe-1/2/0 unit 14 family inet6 address ::10.1.1.14/126
set interfaces fe-1/2/0 unit 14 family mpls
set interfaces lo0 unit 5 family inet6 address ::1.1.1.5/128
set protocols bgp group toPE2 type external
set protocols bgp group toPE2 local-address ::10.1.1.14
set protocols bgp group toPE2 family inet6 unicast
set protocols bgp group toPE2 export send-v6
set protocols bgp group toPE2 peer-as 2
set protocols bgp group toPE2 neighbor ::10.1.1.13
set policy-options policy-statement send-v6 from family inet6
set policy-options policy-statement send-v6 from protocol direct
set policy-options policy-statement send-v6 then accept
set routing-options router-id 1.1.1.5
set routing-options autonomous-system 3
```

Configuring Device PE1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device PE1:

1. Configure the interfaces.


```
[edit interfaces]
user@PE1# set fe-1/2/0 unit 2 family inet6 address ::10.1.1.2/126
user@PE1# set fe-1/2/0 unit 2 family mpls

user@PE1# set fe-1/2/1 unit 5 family inet address 10.1.1.5/30
user@PE1# set fe-1/2/1 unit 5 family inet6
user@PE1# set fe-1/2/1 unit 5 family mpls

user@PE1# set lo0 unit 2 family inet address 1.1.1.2/32
```
2. Configure MPLS on the interfaces.

```
[edit protocols mpls]
user@PE1# set ipv6-tunneling
user@PE1# set interface fe-1/2/0.2
user@PE1# set interface fe-1/2/1.5
```

3. Configure BGP.

```
[edit protocols bgp]
user@PE1# set group toCE1 type external
user@PE1# set group toCE1 local-address ::10.1.1.2
user@PE1# set group toCE1 family inet6 unicast
user@PE1# set group toCE1 export send-bgp6
user@PE1# set group toCE1 peer-as 1
user@PE1# set group toCE1 neighbor ::10.1.1.1

user@PE1# set group toPE2 type internal
user@PE1# set group toPE2 local-address 1.1.1.2
user@PE1# set group toPE2 family inet6 labeled-unicast explicit-null
user@PE1# set group toPE2 export next-hop-self
user@PE1# set group toPE2 export send-v6
user@PE1# set group toPE2 neighbor 1.1.1.4
```

4. Configure OSPF

```
[edit protocols ospf area 0.0.0.0]
user@PE1# set interface fe-1/2/1.5
user@PE1# set interface lo0.2 passive
```

5. Configure a signaling protocol.

```
[edit protocols]
user@PE1# set ldp interface fe-1/2/1.5
```

6. Configure the routing policies.

```
[edit policy-options]
user@PE1# set policy-statement next-hop-self then next-hop self

user@PE1# set policy-statement send-bgp6 from family inet6
user@PE1# set policy-statement send-bgp6 from protocol bgp
user@PE1# set policy-statement send-bgp6 then accept

user@PE1# set policy-statement send-v6 from family inet6
user@PE1# set policy-statement send-v6 from protocol bgp
user@PE1# set policy-statement send-v6 from protocol direct
user@PE1# set policy-statement send-v6 then accept
```

7. Configure the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@PE1# set router-id 1.1.1.2
user@PE1# set autonomous-system 2
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 2 {
    family inet6 {
      address ::10.1.1.2/126;
    }
    family mpls;
  }
}
fe-1/2/1 {
  unit 5 {
    family inet {
      address 10.1.1.5/30;
    }
    family inet6;
    family mpls;
  }
}
lo0 {
  unit 2 {
    family inet {
      address 1.1.1.2/32;
    }
  }
}

user@R1# show policy-options
policy-statement next-hop-self {
  then {
    next-hop self;
  }
}
policy-statement send-bgp6 {
  from {
    family inet6;
    protocol bgp;
  }
  then accept;
}
policy-statement send-v6 {
  from {
    family inet6;
    protocol [ bgp direct ];
  }
  then accept;
}

user@R1# show protocols
mpls {
  ipv6-tunneling;
  interface fe-1/2/0.2;
  interface fe-1/2/1.5;
}
bgp {
  group toCE1 {
    type external;
    local-address ::10.1.1.2;
```

```

family inet6 {
    unicast;
}
export send-bgp6;
peer-as 1;
neighbor ::10.1.1.1;
}
group toPE2 {
    type internal;
    local-address 1.1.1.2;
    family inet6 {
        labeled-unicast {
            explicit-null;
        }
    }
    export [ next-hop-self send-v6 ];
    neighbor 1.1.1.4;
}
}
ospf {
    area 0.0.0.0 {
        interface fe-1/2/1.5;
        interface lo0.2 {
            passive;
        }
    }
}
}
ldp {
    interface fe-1/2/1.5;
}
}

user@R1# show routing-options
router-id 1.1.1.2;
autonomous-system 2;

```

If you are done configuring the device, enter **commit** from configuration mode. Configure the other devices in the topology, as shown in [“CLI Quick Configuration” on page 74](#).

Verification

Confirm that the configuration is working properly.

Verifying That the CE Devices Have Connectivity

Purpose Make sure that the tunnel is operating.

Action From operational mode, enter the **ping** command.

```

user@CE1> ping ::10.1.1.14
PING6(56=40+8+8 bytes) ::10.1.1.1 --> ::10.1.1.14
16 bytes from ::10.1.1.14, icmp_seq=0 hlim=61 time=10.687 ms
16 bytes from ::10.1.1.14, icmp_seq=1 hlim=61 time=9.239 ms
16 bytes from ::10.1.1.14, icmp_seq=2 hlim=61 time=1.842 ms

user@CE3> ping ::10.1.1.1

```

```
PING6(56=40+8+8 bytes) ::10.1.1.14 --> ::10.1.1.1
16 bytes from ::10.1.1.1, icmp_seq=0 hlim=61 time=1.484 ms
16 bytes from ::10.1.1.1, icmp_seq=1 hlim=61 time=1.338 ms
16 bytes from ::10.1.1.1, icmp_seq=2 hlim=61 time=1.351 ms
```

Meaning The IPv6 CE devices can communicate over the core IPv4 network.

Related Documentation

- [Configuring the Ingress Router for MPLS-Signaled LSPs on page 60](#)
- [Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs on page 68](#)
- [Minimum RSVP Configuration on page 471](#)

SRLG Overview

In MPLS traffic engineering, a Shared Risk Link Group (SRLG) is a set of links sharing a common resource, which affects all links in the set if the common resource fails. These links share the same risk of failure and are therefore considered to belong to the same SRLG. For example, links sharing a common fiber are said to be in the same SRLG because a fault with the fiber might cause all links in the group to fail.

An SRLG is represented by a 32-bit number unique within an IGP (OSPFv2 and IS-IS) domain. A link might belong to multiple SRLGs. The SRLG of a path in a label-switched path (LSP) is the set of SRLGs for all the links in the path. When computing the secondary path for an LSP, it is preferable to find a path such that the secondary and primary paths do not have any links in common in case the SRLGs for the primary and secondary paths are disjoint. This ensures that a single point of failure on a particular link does not bring down both the primary and secondary paths in the LSP.

When the SRLG is configured, the device uses the Constrained Shortest Path First (CSPF) algorithm and tries to keep the links used for the primary and secondary paths mutually exclusive. If the primary path goes down, the CSPF algorithm computes the secondary path by trying to avoid links that share any SRLG with the primary path. In addition, when computing the path for a bypass LSP, CSPF tries to avoid links that share any SRLG with the protected links.

When the SRLG is not configured, CSPF only takes into account the costs of the links when computing the secondary path.

Any change in link SRLG information triggers the IGP to send LSP updates for the new link SRLG information. CSPF recomputes the paths during the next round of reoptimization.

Junos OS Release 11.4 and later supports SRLG based on the following RFCs:

- RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*.
- RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*.



NOTE: Currently, the “Fate Sharing” feature continues to be supported with the SRLG feature.

Related Documentation

- [Example: Configuring SRLG on page 81](#)
- [Example: Excluding SRLG Links Completely for the Secondary LSP on page 90](#)
- [Example: Configuring SRLG with Link Protection on page 95](#)
- [Example: Configuring SRLG with Link Protection with the exclude-srlg Option on page 116](#)
- [Computing Backup Paths for LSPs Using Fate Sharing on page 34](#)

Example: Configuring SRLG

This example shows how to configure Shared Risk Link Groups (SRLGs) on a device.

- [Requirements on page 81](#)
- [Overview on page 81](#)
- [Configuration on page 82](#)
- [Verification on page 87](#)

Requirements

This example uses the following hardware and software components:

- Seven routers that can be a combination of M Series, MX Series, or T Series routers
- Junos OS Release 11.4 or later running on all the devices

Overview

Junos OS Release 11.4 and later support SRLG configuration in an IGP (OSPFv2 and IS-IS) domain. In this example, you configure SRLG and associate it with the MPLS interface on a device.

The device uses the SRLG cost parameter for the Constrained Shortest Path First (CSPF) algorithm and tries to keep the links used for the primary and secondary paths mutually exclusive by avoiding links that share any SRLG with the primary path.


To configure the SRLG, you first define the SRLG parameters at the **[edit routing-options srlg srlg-name]** hierarchy level and then associate the SRLG with an MPLS interface at the **[edit mpls interface interface-name]** hierarchy level.

The **srlg srlg-name** statement has the following options:

- **srlg-cost**—Include a cost for the SRLG ranging from 1 through 65535. The cost of the SRLG determines the level of impact this SRLG has on the CSPF algorithm for path computations. The higher the cost, the less likely it is for a secondary path to share the same SRLG as the primary path. By default, the **srlg-cost** is 1.

- **srlg-value**—Include a group ID for the SRLG ranging from 1 through 4294967295.

In this example, PE1 is the ingress router and PE2 is the egress router. P1, P2, and P3, P4, and P5 are transit routers. OSPF is configured on all the routers as the interior gateway protocol (IGP). SRLG is configured on all seven routers. The primary path includes SRLG **srlg-a**. For the standby secondary path, the link P2>PE2 belongs to SRLG **srlg-a**. The effective link metric, with the added **srlg-cost** of 10, becomes 11. Therefore, the computed secondary path is PE1>P3>P4>P5>PE2 with a CSPF link metric of 4.



Configuration

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Router PE1
set interfaces ge-0/0/1 unit 0 family inet address 192.168.12.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.13.1/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 192.168.14.1/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.1/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols rsvp interface ge-0/0/3.0
set protocols mpls optimize-timer 120
set protocols mpls label-switched-path pe1-pe2 to 10.255.0.7
set protocols mpls label-switched-path pe1-pe2 primary via-p1
set protocols mpls label-switched-path pe1-pe2 secondary path2 standby
```



```

set protocols mpls path via-p1 10.255.0.2 strict
set protocols mpls path path2
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/3.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Router P1

```

set interfaces ge-0/0/1 unit 0 family inet address 192.168.12.2/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.27.2/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.2/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0 srlg srlg-a
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Router P2

```

set interfaces ge-0/0/1 unit 0 family inet address 192.168.13.3/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.37.3/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.3/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0 srlg srlg-a
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Router P3

```

set interfaces ge-0/0/1 unit 0 family inet address 192.168.14.4/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.45.4/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.4/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols ospf traffic-engineering

```

```
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0
```

Router P4

```
set interfaces ge-0/0/1 unit 0 family inet address 192.168.45.5/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.56.5/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.5/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0
```

Router P5

```
set interfaces ge-0/0/1 unit 0 family inet address 192.168.56.6/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.67.6/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.6/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0
```

Router PE2

```
set interfaces ge-0/0/1 unit 0 family inet address 192.168.27.7/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.37.7/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 192.168.67.7/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.7/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols rsvp interface ge-0/0/3.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/3.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
```

```
set protocols ospf area 0.0.0.0 interface lo0.0
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure the ingress router PE1:

1. Configure the device interfaces.

```
[edit interfaces]
user@PE1# set ge-0/0/1 unit 0 family inet address 192.168.12.1/24
user@PE1# set ge-0/0/1 unit 0 family mpls
user@PE1# set ge-0/0/2 unit 0 family inet address 192.168.13.1/24
user@PE1# set ge-0/0/2 unit 0 family mpls
user@PE1# set ge-0/0/3 unit 0 family inet address 192.168.14.1/24
user@PE1# set ge-0/0/3 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 10.255.0.1/32
```

2. Configure OSPF on the interfaces.

```
[edit protocols ospf]
user@PE1# set traffic-engineering
user@PE1# set area 0.0.0.0 interface ge-0/0/1.0
user@PE1# set area 0.0.0.0 interface ge-0/0/2.0
user@PE1# set area 0.0.0.0 interface ge-0/0/3.0
user@PE1# set area 0.0.0.0 interface lo0.0
```

3. Configure the SRLG definitions.

```
[edit routing-options]
user@PE1# set srlg srlg-a srlg-value 101
user@PE1# set srlg srlg-a srlg-cost 10
```

4. Configure MPLS and the LSPs.

```
[edit protocols mpls]
user@PE1# set interface ge-0/0/1.0
user@PE1# set interface ge-0/0/2.0
user@PE1# set interface ge-0/0/3.0
user@PE1# set optimize-timer 120
user@PE1# set label-switched-path pe1-pe2 to 10.255.0.7
user@PE1# set label-switched-path pe1-pe2 primary via-p1
user@PE1# set label-switched-path pe1-pe2 secondary path2 standby
user@PE1# set path via-p1 10.255.0.2 strict
user@PE1# set path path2
```

5. Enable RSVP on the interfaces.

```
[edit protocols rsvp]
user@PE1# set interface ge-0/0/1.0
user@PE1# set interface ge-0/0/2.0
user@PE1# set interface ge-0/0/3.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show routing-options**, **show protocols mpls**, and **show protocols rsvp** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show interfaces
interfaces {
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.12.1/24;
      }
      family mpls;
    }
  }
  ge-0/0/2 {
    unit 0 {
      family inet {
        address 192.168.13.1/24;
      }
      family mpls;
    }
  }
  ge-0/0/3 {
    unit 0 {
      family inet {
        address 192.168.14.1/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.0.1/32;
      }
    }
  }
}

user@PE1# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface ge-0/0/1.0;;
  interface ge-0/0/2.0;
  interface ge-0/0/3.0;
  interface lo0.0;
}

user@PE1# show protocols mpls
optimize-timer 120;
label-switched-path pe1-pe2 {
  to 10.255.0.7;
  primary via-p1;
  secondary path2 {
    standby;
  }
}
path via-p1 {
  10.255.0.2 strict;
}
path path2;
```

```

interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

user@PE1# show protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

user@PE1# show routing-options
routing-options {
  srlg {
    srlg-a {
      srlg-value 101;
      srlg-cost 10;
    }
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.



NOTE: Repeat this procedure for every Juniper Networks router in the IGP domain, after modifying the appropriate interface names, addresses, and any other parameters for each router.

Verification

Confirm that the configuration is working properly.

- [Verifying SRLG Definitions on page 87](#)
- [Verify TE-Link SRLG on page 87](#)
- [Verify Standby Secondary Path on page 88](#)

Verifying SRLG Definitions

Purpose Verify SRLG-to-value mappings and SRLG cost.

Action user@PE1> **show mpls srlg**

SRLG	Value	Cost
srlg-a	101	10

Verify TE-Link SRLG

Purpose Verify the traffic engineering link SRLG association.

Action user@PE1> show ted link detail

```
...
10.255.0.2->192.168.27.7-1, Local: 192.168.27.2, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 1, Metric: 1, StaticBW: 1000Mbps, AvailBW: 1000Mbps
  Color: 0 <none>
  SRLGs: srlg-a
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps

...
10.255.0.3->192.168.37.7-1, Local: 192.168.37.3, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 1, StaticBW: 1000Mbps, AvailBW: 1000Mbps
  Color: 0 <none>
  SRLGs: srlg-a
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps

...
```

Meaning Links P1-PE2 and P2-PE2 are associated with SRLG **srlg-a**.

Verify Standby Secondary Path

Purpose Check the SRLG link cost and its impact on the CSPF computation of the standby secondary path link.

Action user@PE1> show mpls lsp ingress extensive
Ingress LSP: 1 sessions

```

10.255.0.7
  From: 10.255.0.1, State: Up, ActiveRoute: 0, LSPname: pe1-pe2
  ActivePath: via-p1 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary via-p1 State: Up
    Priorities: 7 0
    OptimizeTimer: 120
    SmartOptimizeTimer: 180
    SRLG: srlg-a
    Reoptimization in 110 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
192.168.12.2 S 192.168.27.7 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    192.168.12.2 192.168.27.7
    7 Oct 13 15:17:11.310 CSPF: computation result ignored, new path no benefit
    6 Oct 13 15:15:14.959 Selected as active path
    5 Oct 13 15:15:14.958 Record Route: 192.168.12.2 192.168.27.7
    4 Oct 13 15:15:14.954 Up
    3 Oct 13 15:15:14.793 Originate Call
    2 Oct 13 15:15:14.793 CSPF: computation result accepted 192.168.12.2
192.168.27.7
  1 Oct 13 15:14:46.214 CSPF failed: no route toward 10.255.0.2
  Standby path2 State: Up
    Priorities: 7 0
    OptimizeTimer: 120
    SmartOptimizeTimer: 180
    Reoptimization in 115 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 4)
192.168.14.4 S 192.168.45.5 S 192.168.56.6 S 192.168.67.7 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    192.168.14.4 192.168.45.5 192.168.56.6 192.168.67.7
    10 Oct 13 15:17:11.929 Record Route: 192.168.14.4 192.168.45.5 192.168.56.6
192.168.67.7
    9 Oct 13 15:17:11.929 Up
    8 Oct 13 15:17:11.729 Originate Call
    7 Oct 13 15:17:11.729 Clear Call
    6 Oct 13 15:17:11.729 CSPF: computation result accepted 192.168.14.4
192.168.45.5 192.168.56.6 192.168.67.7
    5 Oct 13 15:17:11.729 CSPF: Reroute due to re-optimization
    4 Oct 13 15:15:14.984 Record Route: 192.168.13.3 192.168.37.7
    3 Oct 13 15:15:14.984 Up
    2 Oct 13 15:15:14.830 Originate Call
    1 Oct 13 15:15:14.830 CSPF: computation result accepted 192.168.13.3
192.168.37.7
  Created: Thu Oct 13 15:13:46 2011
  Total 1 displayed, Up 1, Down 0

```

Meaning Check the standby secondary path. The effective link cost for P2>PE2 is 11 (with the added **srlg-cost** of 10). CSPF computes the secondary path as PE1>P3>P4>P5>PE2 with a CSPF link metric of 4.

- Related Documentation**
- [SRLG Overview on page 80](#)
 - [Example: Excluding SRLG Links Completely for the Secondary LSP on page 90](#)
 - [srlg on page 942](#)
 - [srlg-cost on page 943](#)
 - [srlg-value on page 943](#)

Example: Excluding SRLG Links Completely for the Secondary LSP

This example shows how to configure the **exclude-srlg** option to exclude Shared Risk Link Group (SRLG) links for the secondary label-switched path (LSP).

- [Requirements on page 90](#)
- [Overview on page 90](#)
- [Configuration on page 91](#)
- [Verification on page 94](#)

Requirements

This example uses the following hardware and software components:

- M Series, MX Series, or T Series devices
- Junos OS Release 11.4 or later running on all the devices


Overview

For critical links where it is imperative to keep the secondary and primary paths completely disjoint from any common SRLG, you can optionally configure the **exclude-srlg** statement at the **[edit protocols mpls]** or **[edit protocols mpls label-switched-path *path-name*]** hierarchy levels. For logical systems, you configure the **exclude-srlg** statement at the **edit logical-systems protocols mpls[edit logical-systems *logical-system-name* protocols mpls label-switched-path *path-name*]** hierarchy level.

If **exclude-srlg** is configured, the Constrained Shortest Path First (CSPF) algorithm excludes any link belonging to the set of SRLGs in the primary path. If **exclude-srlg** is not configured, and if a link belongs to the set of SRLGs in the primary path, CSPF adds the SRLG cost to the metric, but still accepts the link for computing the path.

In this example, PE1 is the ingress router and PE2 is the egress router. P1, P2, and P3, P4, and P5 are transit routers. OSPF is configured on all the routers as the interior gateway protocol (IGP). SRLG is configured on all seven routers. The primary path includes SRLG **srlg-a**. For the standby secondary path, the link P2>PE2 belongs to SRLG **srlg-a**. Because

`exclude-srlg` is configured, CSPF rejects link P2>PE2 as the link belongs to the SRLG `srlg-a`. Therefore, the computed standby secondary path is PE1>P3>P4>P5>PE2.



Configuration

CLI Quick Configuration To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

Router PE1
set interfaces ge-0/0/1 unit 0 family inet address 192.168.12.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.13.1/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 192.168.14.1/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.1/32
set routing-options srlg srlg-a srlg-value 101
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols rsvp interface ge-0/0/3.0
set protocols mpls optimize-timer 120
set protocols mpls exclude-srlg
set protocols mpls label-switched-path pe1-pe2 to 10.255.0.7
set protocols mpls label-switched-path pe1-pe2 primary via-p1
set protocols mpls label-switched-path pe1-pe2 secondary path2 standby
set protocols mpls path via-p1 10.255.0.2 strict
set protocols mpls path path2
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/3.0
set protocols ospf traffic-engineering
  
```

```

set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

1. Configure the device interfaces.

```

[edit interfaces]
user@PE1# set ge-0/0/1 unit 0 family inet address 192.168.12.1/24
user@PE1# set ge-0/0/1 unit 0 family mpls
user@PE1# set ge-0/0/2 unit 0 family inet address 192.168.13.1/24
user@PE1# set ge-0/0/2 unit 0 family mpls
user@PE1# set ge-0/0/3 unit 0 family inet address 192.168.14.1/24
user@PE1# set ge-0/0/3 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 10.255.0.1/32

```

2. Configure OSPF on the interfaces.

```

[edit protocols ospf]
user@PE1# set traffic-engineering
user@PE1# set area 0.0.0.0 interface ge-0/0/1.0
user@PE1# set area 0.0.0.0 interface ge-0/0/2.0
user@PE1# set area 0.0.0.0 interface ge-0/0/3.0
user@PE1# set area 0.0.0.0 interface lo0.0

```

3. Configure the SRLG definitions.

```

[edit routing-options]
user@PE1# set routing-options srlg srlg-a srlg-value 101

```

4. Configure MPLS and the LSPs.

```

[edit protocols mpls]
user@PE1# set interface ge-0/0/1.0
user@PE1# set interface ge-0/0/2.0
user@PE1# set interface ge-0/0/3.0
user@PE1# set optimize-timer 120
user@PE1# set exclude-srlg
user@PE1# set label-switched-path pe1-pe2 to 10.255.0.7
user@PE1# set label-switched-path pe1-pe2 primary via-p1
user@PE1# set label-switched-path pe1-pe2 secondary path2 standby
user@PE1# set path via-p1 10.255.0.2 strict
user@PE1# set path path2

```

5. Configure the **exclude-srlg** statement to forcibly keep the links for the secondary path completely disjoint from the primary LSP path.

```

user@PE1 set protocols mpls exclude-srlg

```

6. Enable RSVP on the interfaces.

```

[edit protocols rsvp]
user@PE1# set interface ge-0/0/1.0
user@PE1# set interface ge-0/0/2.0
user@PE1# set interface ge-0/0/3.0

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show routing-options**, **show protocols mpls**, and **show protocols rsvp** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show interfaces
interfaces {
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.12.1/24;
      }
      family mpls;
    }
  }
  ge-0/0/2 {
    unit 0 {
      family inet {
        address 192.168.13.1/24;
      }
      family mpls;
    }
  }
  ge-0/0/3 {
    unit 0 {
      family inet {
        address 192.168.14.1/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.0.1/32;
      }
    }
  }
}
```

```
user@PE1# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface ge-0/0/1.0;;
  interface ge-0/0/2.0;
  interface ge-0/0/3.0;
  interface lo0.0;
}
```

```
user@PE1# show protocols mpls
optimize-timer 120;
label-switched-path pe1-pe2 {
  to 10.255.0.7;
  primary via-p1;
  secondary path2 {
    standby;
  }
}
```

```
    }  
  }  
  path via-p1 {  
    10.255.0.2 strict;  
  }  
  path path2;  
  interface ge-0/0/1.0;  
  interface ge-0/0/2.0;  
  interface ge-0/0/3.0;  
  
user@PE1# show protocols rsvp  
interface ge-0/0/1.0;  
interface ge-0/0/2.0;  
interface ge-0/0/3.0;  
  
user@PE1# show routing-options  
routing-options {  
  srlg {  
    srlg-a srlg-value 101;  
  }  
}
```

If you are done configuring the device, enter **commit** from configuration mode.



NOTE: Repeat this procedure for every Juniper Networks router in the IGP domain, after modifying the appropriate interface names, addresses, and any other parameters for each router.

Verification

Confirm that the configuration is working properly.

Verifying the Secondary Path Link for the LSP

Purpose Verify that the link for the secondary path is completely disjoint from the primary path.

Action user@PE1> show mpls lsp detail
Ingress LSP: 1 sessions

```

10.255.0.7
  From: 10.255.0.1, State: Up, ActiveRoute: 0, LSPname: pe1-pe2
  ActivePath: via-p1 (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary via-p1 State: Up
    Priorities: 7 0
    OptimizeTimer: 120
    SmartOptimizeTimer: 180
    SRLG: srlg-a
    Reoptimization in 77 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
192.168.12.2 S 192.168.27.7 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    192.168.12.2 192.168.27.7
  Standby path2 State: Up
    Priorities: 7 0
    OptimizeTimer: 120
    SmartOptimizeTimer: 180
    Reoptimization in 106 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 4)
192.168.14.4 S 192.168.45.5 S 192.168.56.6 S 192.168.67.7 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    192.168.14.4 192.168.45.5 192.168.56.6 192.168.67.7
Total 1 displayed, Up 1, Down 0

Link P1->PE2: SRLG srlg-a
Link P2->PE2: SRLG srlg-a

Primary path: PE1-P1-PE2 (CSPF metric: 2)
Standby secondary: PE1-P3-P4-P5-PE2 (CSPF metric: 4)

```

Meaning Primary path includes SRLG **srlg-a**. For the standby secondary path, the link P2>PE2 belongs to SRLG **srlg-a**. CSPF rejects link P2>PE2 because the link belongs to the SRLG **srlg-a**.

- Related Documentation**
- [SRLG Overview on page 80](#)
 - [Example: Configuring SRLG on page 81](#)
 - [exclude-srlg on page 863](#)

Example: Configuring SRLG with Link Protection

This example shows how to configure SRLG with link protection without the **exclude-srlg** option.

- [Requirements on page 96](#)
- [Overview on page 96](#)

- [Configuration on page 96](#)
- [Verification on page 114](#)

Requirements

This example uses the following hardware and software components:


- M Series, MX Series, or T Series devices
- Junos OS Release 11.4 or later running on all the devices

Overview

In this example, PE1 is the ingress router and PE2 is the egress router. P1, P2, and P3, P4, and P5 are transit routers. OSPF is configured on all the routers as the interior gateway protocol (IGP). SRLG is configured on all seven routers. The link P1>PE2 (primary path) and the link P2>PE2 belong to SRLG srlg-a.

You configure link protection for the interface P1>PE2 by including the **link-protection** statement.

When SRLG srlg-a is configured on the link P1>PE2 and P2>PE2, the bypass takes the longer path P1>P4>P5>PE2, not selecting the link P2>PE2 because of the added SRLG cost for srlg-a.



Configuration

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router PE1

```

set interfaces ge-0/0/1 unit 0 family inet address 192.168.12.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.13.1/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 192.168.14.1/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.1/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols rsvp interface ge-0/0/3.0
set protocols mpls optimize-timer 120
set protocols mpls label-switched-path pe1-pe2 to 10.255.0.7
set protocols mpls label-switched-path pe1-pe2 link-protection
set protocols mpls label-switched-path pe1-pe2 primary via-p1
set protocols mpls label-switched-path pe1-pe2 secondary path2 standby
set protocols mpls path via-p1 10.255.0.2 strict
set protocols mpls path path2
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/3.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Router P1

```

set interfaces ge-0/0/1 unit 0 family inet address 192.168.12.2/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.27.2/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 192.168.23.2/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces ge-0/0/4 unit 0 family inet address 192.168.25.2/24
set interfaces ge-0/0/4 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.2/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0 link-protection
set protocols rsvp interface ge-0/0/3.0
set protocols rsvp interface ge-0/0/4.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0 srlg srlg-a
set protocols mpls interface ge-0/0/3.0
set protocols mpls interface ge-0/0/4.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface ge-0/0/4.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Router P2 set interfaces ge-0/0/1 unit 0 family inet address 192.168.13.3/24

```
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.37.3/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 192.168.23.3/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.3/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols rsvp interface ge-0/0/3.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0 srlg srlg-a
set protocols mpls interface ge-0/0/3.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0
```

Router P3

```
set interfaces ge-0/0/1 unit 0 family inet address 192.168.14.4/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.45.4/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.4/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0
```

Router P4

```
set interfaces ge-0/0/1 unit 0 family inet address 192.168.45.5/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.56.5/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 192.168.25.5/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.5/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols rsvp interface ge-0/0/3.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/3.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
```



```
set protocols ospf area 0.0.0.0 interface lo0.0
```

Router P5

```
set interfaces ge-0/0/1 unit 0 family inet address 192.168.56.6/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.67.6/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.6/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0
```

Router PE2

```
set interfaces ge-0/0/1 unit 0 family inet address 192.168.27.7/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.37.7/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 192.168.67.7/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.7/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols rsvp interface ge-0/0/3.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/3.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0
```

Configuring Device PE1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure the ingress router PE1:

1. Configure the device interfaces.

```
[edit interfaces]
user@PE1# set ge-0/0/1 unit 0 family inet address 192.168.12.1/24
user@PE1# set ge-0/0/1 unit 0 family mpls
user@PE1# set ge-0/0/2 unit 0 family inet address 192.168.13.1/24
user@PE1# set ge-0/0/2 unit 0 family mpls
user@PE1# set ge-0/0/3 unit 0 family inet address 192.168.14.1/24
user@PE1# set ge-0/0/3 unit 0 family mpls
```

```
user@PE1# set lo0 unit 0 family inet address 10.255.0.1/32
```

2. Configure OSPF on the interfaces.

```
[edit protocols ospf]
user@PE1# set traffic-engineering
user@PE1# set area 0.0.0.0 interface ge-0/0/1.0
user@PE1# set area 0.0.0.0 interface ge-0/0/2.0
user@PE1# set area 0.0.0.0 interface ge-0/0/3.0
user@PE1# set area 0.0.0.0 interface lo0.0
```

3. Configure the SRLG definitions.

```
[edit routing-options]
user@PE1# set srlg srlg-a srlg-value 101
user@PE1# set srlg srlg-a srlg-cost 10
```

4. Configure MPLS and the LSPs and configure link protection for the **pe1-pe2** LSP.

```
[edit protocols mpls]
user@PE1# set interface ge-0/0/1.0
user@PE1# set interface ge-0/0/2.0
user@PE1# set interface ge-0/0/3.0
user@PE1# set optimize-timer 120
user@PE1# set label-switched-path pe1-pe2 to 10.255.0.7
user@PE1# set protocols mpls label-switched-path pe1-pe2 link-protection
user@PE1# set label-switched-path pe1-pe2 primary via-p1
user@PE1# set label-switched-path pe1-pe2 secondary path2 standby
user@PE1# set path via-p1 10.255.0.2 strict
user@PE1# set path path2
```

5. Enable RSVP on the interfaces.

```
[edit protocols rsvp]
user@PE1# set interface ge-0/0/1.0
user@PE1# set interface ge-0/0/2.0
user@PE1# set interface ge-0/0/3.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show routing-options**, **show protocols mpls**, and **show protocols rsvp** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.12.1/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 192.168.13.1/24;
    }
    family mpls;
  }
}
```

```

    }
  }
  ge-0/0/3 {
    unit 0 {
      family inet {
        address 192.168.14.1/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.0.1/32;
      }
    }
  }
}

user@PE1# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface ge-0/0/1.0;
  interface ge-0/0/2.0;
  interface ge-0/0/3.0;
  interface lo0.0;
}

user@PE1# show protocols mpls
optimize-timer 120;
label-switched-path pe1-pe2 {
  to 10.255.0.7;
  link-protection;
  primary via-p1;
  secondary path2 {
    standby;
  }
}
path via-p1 {
  10.255.0.2 strict;
}
path path2;
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

user@PE1# show protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

user@PE1# show routing-options
srlg {
  srlg-a {
    srlg-value 101;
    srlg-cost 10;
  }
}

```

```
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device P1

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure device P1:

1. Configure the device interfaces.

[edit interfaces]

```
user@P1# set ge-0/0/1 unit 0 family inet address 192.168.12.2/24
user@P1# set ge-0/0/1 unit 0 family mpls
user@P1# set ge-0/0/2 unit 0 family inet address 192.168.27.2/24
user@P1# set ge-0/0/2 unit 0 family mpls
user@P1# set ge-0/0/3 unit 0 family inet address 192.168.23.2/24
user@P1# set ge-0/0/3 unit 0 family mpls
user@P1# set ge-0/0/4 unit 0 family inet address 192.168.25.2/24
user@P1# set ge-0/0/4 unit 0 family mpls
user@P1# set lo0 unit 0 family inet address 10.255.0.2/32
```

2. Configure OSPF on the interfaces.

[edit protocols ospf]

```
user@P1# set traffic-engineering
user@P1# set area 0.0.0.0 interface ge-0/0/1.0
user@P1# set area 0.0.0.0 interface ge-0/0/2.0
user@P1# set area 0.0.0.0 interface ge-0/0/3.0
user@P1# set area 0.0.0.0 interface ge-0/0/4.0
user@P1# set area 0.0.0.0 interface lo0.0
```

3. Configure the SRLG definitions.

[edit routing-options]

```
user@P1# set srlg srlg-a srlg-value 101
user@P1# set srlg srlg-a srlg-cost 10
```

4. Configure MPLS on the interfaces and associate the SRLG **srlg-a** with interface **ge-0/0/2.0** for the P1>PE2 link.

[edit protocols mpls]

```
user@P1# set interface ge-0/0/1.0
user@P1# set interface ge-0/0/2.0 srlg srlg-a
user@P1# set interface ge-0/0/3.0
user@P1# set interface ge-0/0/4.0
```

5. Enable RSVP on the interfaces and configure **link-protection** for interface **ge-0/0/2.0**.

[edit protocols rsvp]

```
user@P1# set interface ge-0/0/1.0
user@P1# set interface ge-0/0/2.0 link-protection
user@P1# set interface ge-0/0/3.0
user@P1# set interface ge-0/0/4.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show protocols mpls**, **show protocols RSVP**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@P1# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.12.2/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 192.168.27.2/24;
    }
    family mpls;
  }
}
ge-0/0/3 {
  unit 0 {
    family inet {
      address 192.168.23.2/24;
    }
    family mpls;
  }
}
ge-0/0/4 {
  unit 0 {
    family inet {
      address 192.168.25.2/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.0.2/32;
    }
  }
}

user@P1# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface ge-0/0/1.0;
  interface ge-0/0/2.0;
  interface ge-0/0/3.0;
  interface ge-0/0/4.0;
  interface lo0.0;
}

```

```
user@P1# show protocols mpls
interface ge-0/0/1.0;
interface ge-0/0/2.0 {
  srlg srlg-a;
}
interface ge-0/0/3.0;
interface ge-0/0/4.0;

user@P1# show protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0 {
  link-protection;
}
interface ge-0/0/3.0;
interface ge-0/0/4.0;

user@P1# show routing-options
srlg {
  srlg-a {
    srlg-value 101;
    srlg-cost 10;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device P2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure P2:

1. Configure the device interfaces.

[edit interfaces]

```
user@P2# set ge-0/0/1 unit 0 family inet address 192.168.13.3/24
user@P2# set ge-0/0/1 unit 0 family mpls
user@P2# set ge-0/0/2 unit 0 family inet address 192.168.37.3/24
user@P2# set ge-0/0/2 unit 0 family mpls
user@P2# set ge-0/0/3 unit 0 family inet address 192.168.23.3/24
user@P2# set ge-0/0/3 unit 0 family mpls
user@P2# set lo0 unit 0 family inet address 10.255.0.3/32
```

2. Configure OSPF on the interfaces.

[edit protocols ospf]

```
user@P2# set traffic-engineering
user@P2# set area 0.0.0.0 interface ge-0/0/1.0
user@P2# set area 0.0.0.0 interface ge-0/0/2.0
user@P2# set area 0.0.0.0 interface ge-0/0/3.0
user@P2# set area 0.0.0.0 interface lo0.0
```

3. Configure the SRLG definitions.

[edit routing-options]

```
user@P2# set srlg srlg-a srlg-value 101
user@P2# set srlg srlg-a srlg-cost 10
```

4. Configure MPLS on the interfaces and associate the SRLG **srlg-a** with interface **ge-0/0/2.0** for the P2>PE2 link.

```
[edit protocols mpls]
user@P2# set interface ge-0/0/1.0
user@P2# set interface ge-0/0/2.0 srlg srlg-a
user@P2# set interface ge-0/0/3.0
```

5. Enable RSVP on the interfaces.

```
[edit protocols rsvp]
user@P2# set interface ge-0/0/1.0
user@P2# set interface ge-0/0/2.0
user@P2# set interface ge-0/0/3.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show protocols mpls**, **show protocols rsvp**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P2# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.13.3/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 192.168.37.3/24;
    }
    family mpls;
  }
}
ge-0/0/3 {
  unit 0 {
    family inet {
      address 192.168.23.3/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.0.3/32;
    }
  }
}
}
```

```
user@P2# show protocols ospf
traffic-engineering;
```

```

area 0.0.0.0 {
    interface ge-0/0/1.0;
    interface ge-0/0/2.0;
    interface ge-0/0/3.0;
    interface lo0.0;
}

user@P2# show protocols mpls
interface ge-0/0/1.0;
interface ge-0/0/2.0 {
    srlg srlg-a;
}
interface ge-0/0/3.0;
}

user@P2# show protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

user@P2# show routing-options
srlg {
    srlg-a {
        srlg-value 101;
        srlg-cost 10;
    }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device P3

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure P3:

1. Configure the device interfaces.

[edit interfaces]

```

user@P3# set ge-0/0/1 unit 0 family inet address 192.168.14.4/24
user@P3# set ge-0/0/1 unit 0 family mpls
user@P3# set ge-0/0/2 unit 0 family inet address 192.168.45.4/24
user@P3# set ge-0/0/2 unit 0 family mpls
user@P3# set lo0 unit 0 family inet address 10.255.0.4/32

```

2. Configure OSPF on the interfaces.

[edit protocols ospf]

```

user@P3# set traffic-engineering
user@P3# set area 0.0.0.0 interface ge-0/0/1.0
user@P3# set area 0.0.0.0 interface ge-0/0/2.0
user@P3# set area 0.0.0.0 interface lo0.0

```

3. Configure the SRLG definitions.

[edit routing-options]

```

user@P3# set srlg srlg-a srlg-value 101
user@P3# set srlg srlg-a srlg-cost 10

```


4. Configure MPLS on the interfaces.

```
[edit protocols mpls]
user@P3# set interface ge-0/0/1.0
user@P3# set interface ge-0/0/2.0
```

5. Enable RSVP on the interfaces.

```
[edit protocols rsvp]
user@P3# set interface ge-0/0/1.0
user@P3# set interface ge-0/0/2.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show protocols mpls**, **show protocols rsvp**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P3# show interfaces
interfaces {
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.14.4/24;
      }
      family mpls;
    }
  }
  ge-0/0/2 {
    unit 0 {
      family inet {
        address 192.168.45.4/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.0.4/32;
      }
    }
  }
}

user@P3# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface ge-0/0/1.0;
  interface ge-0/0/2.0;
  interface lo0.0;
}

user@P3# show protocols mpls
interface ge-0/0/1.0;
interface ge-0/0/2.0;

user@P3# show protocols rsvp
```

```
interface ge-0/0/1.0;
interface ge-0/0/2.0;

user@P3# show routing-options
srlg {
  srlg-a {
    srlg-value 101;
    srlg-cost 10;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device P4

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure P4:

1. Configure the device interfaces.

[edit interfaces]

```
user@P4# set ge-0/0/1 unit 0 family inet address 192.168.45.5/24
user@P4# set ge-0/0/1 unit 0 family mpls
user@P4# set ge-0/0/2 unit 0 family inet address 192.168.56.5/24
user@P4# set ge-0/0/2 unit 0 family mpls
user@P4# set ge-0/0/3 unit 0 family inet address 192.168.25.5/24
user@P4# set ge-0/0/3 unit 0 family mpls
user@P4# set lo0 unit 0 family inet address 10.255.0.5/32
```

2. Configure OSPF on the interfaces.

[edit protocols ospf]

```
user@P4# set traffic-engineering
user@P4# set area 0.0.0.0 interface ge-0/0/1.0
user@P4# set area 0.0.0.0 interface ge-0/0/2.0
user@P4# set area 0.0.0.0 interface ge-0/0/3.0
user@P4# set area 0.0.0.0 interface lo0.0
```

3. Configure the SRLG definitions.

[edit routing-options]

```
user@P4# set srlg srlg-a srlg-value 101
user@P4# set srlg srlg-a srlg-cost 10
```

4. Configure MPLS on the interfaces.

[edit protocols mpls]

```
user@P4# set interface ge-0/0/1.0
user@P4# set interface ge-0/0/2.0
user@P4# set interface ge-0/0/3.0
```

5. Enable RSVP on the interfaces.

[edit protocols rsvp]

```
user@P4# set interface ge-0/0/1.0
user@P4# set interface ge-0/0/2.0
user@P4# set interface ge-0/0/3.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show protocols mpls**, **show protocols rsvp**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@P4# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.45.5/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 192.168.56.5/24;
    }
    family mpls;
  }
}
ge-0/0/3 {
  unit 0 {
    family inet {
      address 192.168.25.5/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.0.5/32;
    }
  }
}

```

```

user@P4# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface ge-0/0/1.0;
  interface ge-0/0/2.0;
  interface ge-0/0/3.0;
  interface lo0.0;
}

```

```

user@P4# show protocols mpls
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

```

```

user@P4# show protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

```

```
user@P4# show routing-options
srlg {
  srlg-a {
    srlg-value 101;
    srlg-cost 10;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device P5

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure P5:

1. Configure the device interfaces.

```
[edit interfaces]
user@P5# set ge-0/0/1 unit 0 family inet address 192.168.56.6/24
user@P5# set ge-0/0/1 unit 0 family mpls
user@P5# set ge-0/0/2 unit 0 family inet address 192.168.67.6/24
user@P5# set ge-0/0/2 unit 0 family mpls
user@P5# set lo0 unit 0 family inet address 10.255.0.6/32
```

2. Configure OSPF on the interfaces.

```
[edit protocols ospf]
user@P5# set traffic-engineering
user@P5# set area 0.0.0.0 interface ge-0/0/1.0
user@P5# set area 0.0.0.0 interface ge-0/0/2.0
user@P5# set area 0.0.0.0 interface lo0.0
```

3. Configure the SRLG definitions.

```
[edit routing-options]
user@P5# set srlg srlg-a srlg-value 101
user@P5# set srlg srlg-a srlg-cost 10
```

4. Configure MPLS on the interfaces.

```
[edit protocols mpls]
user@P5# set interface ge-0/0/1.0
user@P5# set interface ge-0/0/2.0
```

5. Enable RSVP on the interfaces.

```
[edit protocols rsvp]
user@P5# set interface ge-0/0/1.0
user@P5# set interface ge-0/0/2.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show protocols mpls**, **show protocols rsvp**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P5# show interfaces
```

```

ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.56.6/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 192.168.67.6/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.0.6/32;
    }
  }
}

```

```

user@P5# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface ge-0/0/1.0;
  interface ge-0/0/2.0;
  interface lo0.0;
}

```

```

user@P5# show protocols mpls
interface ge-0/0/1.0;
interface ge-0/0/2.0;

```

```

user@P5# show protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0;

```

```

user@P5# show routing-options
srlg {
  srlg-a {
    srlg-value 101;
    srlg-cost 10;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device PE2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure PE2:

1. Configure the device interfaces.

```
[edit interfaces]
user@PE2# set ge-0/0/1 unit 0 family inet address 192.168.27.7/24
user@PE2# set ge-0/0/1 unit 0 family mpls
user@PE2# set ge-0/0/2 unit 0 family inet address 192.168.37.7/24
user@PE2# set ge-0/0/2 unit 0 family mpls
user@PE2# set ge-0/0/3 unit 0 family inet address 192.168.67.7/24
user@PE2# set ge-0/0/3 unit 0 family mpls
user@PE2# set lo0 unit 0 family inet address 10.255.0.7/32
```

2. Configure OSPF on the interfaces.

```
[edit protocols ospf]
user@PE2# set traffic-engineering
user@PE2# set area 0.0.0.0 interface ge-0/0/1.0
user@PE2# set area 0.0.0.0 interface ge-0/0/2.0
user@PE2# set area 0.0.0.0 interface ge-0/0/3.0
user@PE2# set area 0.0.0.0 interface lo0.0
```

3. Configure the SRLG definitions.

```
[edit routing-options]
user@PE2# set srlg srlg-a srlg-value 101
user@PE2# set srlg srlg-a srlg-cost 10
```

4. Configure MPLS on the interfaces.

```
[edit protocols mpls]
user@PE2# set interface ge-0/0/1.0
user@PE2# set interface ge-0/0/2.0
user@PE2# set interface ge-0/0/3.0
```

5. Enable RSVP on the interfaces.

```
[edit protocols rsvp]
user@PE2# set interface ge-0/0/1.0
user@PE2# set interface ge-0/0/2.0
user@PE2# set interface ge-0/0/3.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show protocols mpls**, **show protocols rsvp**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE2# show interfaces
interfaces {
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.27.7/24;
```

```

    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 192.168.37.7/24;
    }
    family mpls;
  }
}
ge-0/0/3 {
  unit 0 {
    family inet {
      address 192.168.67.7/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.0.7/32;
    }
  }
}
}

```

```

user@PE2# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface ge-0/0/1.0;
  interface ge-0/0/2.0;
  interface ge-0/0/3.0;
  interface lo0.0;
}

```

```

user@PE2# show protocols mpls
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

```

```

user@PE2# show protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

```

```

user@PE2# show routing-options
srlg {
  srlg-a {
    srlg-value 101;
    srlg-cost 10;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying the SRLG Cost Is Added to the TE Link

Purpose Verify that the SRLG cost is added to the TE link if it belongs to the SRLG of the protected link. Issue the **show ted link detail** and **show rsvp session extensive bypass** commands on device P1.

Action user@P1> show ted link detail

```
...
10.255.0.2->192.168.27.7-1, Local: 192.168.27.2, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 1, StaticBW: 1000Mbps, AvailBW: 1000Mbps
  Color: 0 <none>
  SRLGs: srlg-a
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
[...]
10.255.0.3->192.168.37.7-1, Local: 192.168.37.3, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 1, StaticBW: 1000Mbps, AvailBW: 1000Mbps
  Color: 0 <none>
  SRLGs: srlg-a
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
  localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
...
```

user@P1> show rsvp session extensive bypass

```
Ingress RSVP: 1 sessions

10.255.0.7
  From: 10.255.0.2, LSPstate: Up, ActiveRoute: 0
  LSName: Bypass->192.168.27.7
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299776
  Resv style: 1 SE, Label in: -, Label out: 299776
  Time left: -, Since: Fri Oct 21 13:19:21 2011
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 52081 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 1
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 192.168.25.5 (ge-0/0/4.0) 26 pkts
  RESV rcvfrom: 192.168.25.5 (ge-0/0/4.0) 26 pkts
  Explct route: 192.168.25.5 192.168.56.6 192.168.67.7
  Record route: <self> 192.168.25.5 192.168.56.6 192.168.67.7
Total 1 displayed, Up 1, Down 0
```

Meaning The shortest path for the bypass protecting the link P1->PE2 would have been P1->P2->PE2. Because the links P1>PE2 and P2>PE2 both belong to SRLG **srlg-a**, the SRLG cost of 10 for **srlg-a** is added to the metric for the link P2>PE2. This makes the metric for the link P2>PE2 too high to be selected for the shortest path. Therefore, the CSPF result for the computed path for the bypass becomes P1>P4>P5>PE2.

Related Documentation

- [SRLG Overview on page 80](#)
- [Example: Configuring SRLG on page 81](#)
- [Example: Configuring SRLG with Link Protection with the exclude-srlg Option on page 116](#)

Example: Configuring SRLG with Link Protection with the `exclude-srlg` Option

This example shows how to configure SRLG with link protection with the **`exclude-srlg`** option.

- [Requirements on page 116](#)
- [Overview on page 116](#)
- [Configuration on page 117](#)
- [Verification on page 134](#)

Requirements

This example uses the following hardware and software components:


- M Series, MX Series, or T Series devices
- Junos OS Release 11.4 or later running on all the devices

Overview

In this example, PE1 is the ingress router and PE2 is the egress router. P1, P2, and P3, P4, and P5 are transit routers. OSPF is configured on all the routers as the interior gateway protocol (IGP). SRLG is configured on all seven routers. The link P1>PE2 (primary path) and the link P2>PE2 belong to SRLG **`srlg-a`**.

You configure link protection for the interface P1>PE2 by including the **`link-protection`** statement along with the **`exclude-srlg`** option. This makes the bypass LSP and the protected link completely disjoint in any SRLG.

When SRLG **srlg-a** is configured on the link P1>PE2 and P2>PE2, the link P2>PE2 is rejected for CSPF consideration due to the **exclude-srlg** configuration. Therefore, the computed path for the bypass becomes P1>P4>P5>PE2.



Configuration

CLI Quick Configuration

To quickly configure this section of the example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Router PE1

```
set interfaces ge-0/0/1 unit 0 family inet address 192.168.12.1/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.13.1/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 192.168.14.1/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.1/32
set routing-options srlg srlg-a srlg-value 101
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols rsvp interface ge-0/0/3.0
set protocols mpls optimize-timer 120
set protocols mpls label-switched-path pe1-pe2 to 10.255.0.7
set protocols mpls label-switched-path pe1-pe2 link-protection
set protocols mpls label-switched-path pe1-pe2 primary via-p1
set protocols mpls label-switched-path pe1-pe2 secondary path2 standby
set protocols mpls path via-p1 10.255.0.2 strict
set protocols mpls path path2
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/3.0
```

```
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0
```

```
Router P1  set interfaces ge-0/0/1 unit 0 family inet address 192.168.12.2/24
            set interfaces ge-0/0/1 unit 0 family mpls
            set interfaces ge-0/0/2 unit 0 family inet address 192.168.27.2/24
            set interfaces ge-0/0/2 unit 0 family mpls
            set interfaces ge-0/0/3 unit 0 family inet address 192.168.23.2/24
            set interfaces ge-0/0/3 unit 0 family mpls
            set interfaces ge-0/0/4 unit 0 family inet address 192.168.25.2/24
            set interfaces ge-0/0/4 unit 0 family mpls
            set interfaces lo0 unit 0 family inet address 10.255.0.2/32
            set routing-options srlg srlg-a srlg-value 101
            set routing-options srlg srlg-a srlg-cost 10
            set protocols rsvp interface ge-0/0/1.0
            set protocols rsvp interface ge-0/0/2.0 link-protection exclude-srlg
            set protocols mpls interface ge-0/0/1.0
            set protocols mpls interface ge-0/0/2.0 srlg srlg-a
            set protocols mpls interface ge-0/0/3.0
            set protocols mpls interface ge-0/0/4.0
            set protocols ospf traffic-engineering
            set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
            set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
            set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
            set protocols ospf area 0.0.0.0 interface ge-0/0/4.0
            set protocols ospf area 0.0.0.0 interface lo0.0
```

```
Router P2  set interfaces ge-0/0/1 unit 0 family inet address 192.168.13.3/24
            set interfaces ge-0/0/1 unit 0 family mpls
            set interfaces ge-0/0/2 unit 0 family inet address 192.168.37.3/24
            set interfaces ge-0/0/2 unit 0 family mpls
            set interfaces ge-0/0/3 unit 0 family inet address 192.168.23.3/24
            set interfaces ge-0/0/3 unit 0 family mpls
            set interfaces lo0 unit 0 family inet address 10.255.0.3/32
            set routing-options srlg srlg-a srlg-value 101
            set routing-options srlg srlg-a srlg-cost 10
            set protocols rsvp interface ge-0/0/1.0
            set protocols rsvp interface ge-0/0/2.0
            set protocols rsvp interface ge-0/0/3.0
            set protocols mpls interface ge-0/0/1.0
            set protocols mpls interface ge-0/0/2.0 srlg srlg-a
            set protocols mpls interface ge-0/0/3.0
            set protocols ospf traffic-engineering
            set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
            set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
            set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
            set protocols ospf area 0.0.0.0 interface lo0.0
```

```
Router P3  set interfaces ge-0/0/1 unit 0 family inet address 192.168.14.4/24
            set interfaces ge-0/0/1 unit 0 family mpls
            set interfaces ge-0/0/2 unit 0 family inet address 192.168.45.4/24
            set interfaces ge-0/0/2 unit 0 family mpls
```

```

set interfaces lo0 unit 0 family inet address 10.255.0.4/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Router P4

```

set interfaces ge-0/0/1 unit 0 family inet address 192.168.45.5/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.56.5/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 192.168.25.5/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.5/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols rsvp interface ge-0/0/3.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/3.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Router P5

```

set interfaces ge-0/0/1 unit 0 family inet address 192.168.56.6/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.67.6/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.0.6/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Router PE2

```

set interfaces ge-0/0/1 unit 0 family inet address 192.168.27.7/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 192.168.37.7/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces ge-0/0/3 unit 0 family inet address 192.168.67.7/24
set interfaces ge-0/0/3 unit 0 family mpls

```

```

set interfaces lo0 unit 0 family inet address 10.255.0.7/32
set routing-options srlg srlg-a srlg-value 101
set routing-options srlg srlg-a srlg-cost 10
set protocols rsvp interface ge-0/0/1.0
set protocols rsvp interface ge-0/0/2.0
set protocols rsvp interface ge-0/0/3.0
set protocols mpls interface ge-0/0/1.0
set protocols mpls interface ge-0/0/2.0
set protocols mpls interface ge-0/0/3.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Configuring Device PE1

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure the ingress router PE1:

1. Configure the device interfaces.

[edit interfaces]

```

user@PE1# set ge-0/0/1 unit 0 family inet address 192.168.12.1/24
user@PE1# set ge-0/0/1 unit 0 family mpls
user@PE1# set ge-0/0/2 unit 0 family inet address 192.168.13.1/24
user@PE1# set ge-0/0/2 unit 0 family mpls
user@PE1# set ge-0/0/3 unit 0 family inet address 192.168.14.1/24
user@PE1# set ge-0/0/3 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 10.255.0.1/32

```

2. Configure OSPF on the interfaces.

[edit protocols ospf]

```

user@PE1# set traffic-engineering
user@PE1# set area 0.0.0.0 interface ge-0/0/1.0
user@PE1# set area 0.0.0.0 interface ge-0/0/2.0
user@PE1# set area 0.0.0.0 interface ge-0/0/3.0
user@PE1# set area 0.0.0.0 interface lo0.0

```

3. Configure the SRLG definitions.

[edit routing-options]

```

user@PE1# set routing-options srlg srlg-a srlg-value 101
user@PE1# set routing-options srlg srlg-a srlg-cost 10

```

4. Configure MPLS and the LSPs and configure link protection for the **pe1-pe2** LSP.

[edit protocols mpls]

```

user@PE1# set interface ge-0/0/1.0
user@PE1# set interface ge-0/0/2.0
user@PE1# set interface ge-0/0/3.0
user@PE1# set optimize-timer 120
user@PE1# set label-switched-path pe1-pe2 to 10.255.0.7
user@PE1# set protocols mpls label-switched-path pe1-pe2 link-protection
user@PE1# set label-switched-path pe1-pe2 primary via-pl

```

```

user@PE1# set label-switched-path pe1-pe2 secondary path2 standby
user@PE1# set path via-p1 10.255.0.2 strict
user@PE1# set path path2

```

5. Enable RSVP on the interfaces.

```

[edit protocols rsvp]
user@PE1# set interface ge-0/0/1.0
user@PE1# set interface ge-0/0/2.0
user@PE1# set interface ge-0/0/3.0

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show routing-options**, **show protocols mpls**, and **show protocols rsvp** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PE1# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.12.1/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 192.168.13.1/24;
    }
    family mpls;
  }
}
ge-0/0/3 {
  unit 0 {
    family inet {
      address 192.168.14.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.0.1/32;
    }
  }
}
}

user@PE1# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface ge-0/0/1.0;
  interface ge-0/0/2.0;
  interface ge-0/0/3.0;
}

```

```
interface lo0.0;
}

user@PE1# show protocols mpls
optimize-timer 120;
label-switched-path pe1-pe2 {
  to 10.255.0.7;
  link-protection;
  primary via-p1;
  secondary path2 {
    standby;
  }
}
path via-p1 {
  10.255.0.2 strict;
}
path path2;
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

user@PE1# show protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

user@PE1# show routing-options
srlg {
  srlg-a {
    srlg-value 101;
    srlg-cost 10;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device P1

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure device P1:

1. Configure the device interfaces.

[edit interfaces]

```
user@P1# set ge-0/0/1 unit 0 family inet address 192.168.12.2/24
user@P1# set ge-0/0/1 unit 0 family mpls
user@P1# set ge-0/0/2 unit 0 family inet address 192.168.27.2/24
user@P1# set ge-0/0/2 unit 0 family mpls
user@P1# set ge-0/0/3 unit 0 family inet address 192.168.23.2/24
user@P1# set ge-0/0/3 unit 0 family mpls
user@P1# set ge-0/0/4 unit 0 family inet address 192.168.25.2/24
user@P1# set ge-0/0/4 unit 0 family mpls
user@P1# set lo0 unit 0 family inet address 10.255.0.2/32
```

2. Configure OSPF on the interfaces.


```
[edit protocols ospf]
user@P1# set traffic-engineering
user@P1# set area 0.0.0.0 interface ge-0/0/1.0
user@P1# set area 0.0.0.0 interface ge-0/0/2.0
user@P1# set area 0.0.0.0 interface ge-0/0/3.0
user@P1# set area 0.0.0.0 interface ge-0/0/4.0
user@P1# set area 0.0.0.0 interface lo0.0
```

3. Configure the SRLG definitions.

```
[edit routing-options]
user@P1# set routing-options srlg srlg-a srlg-value 101
user@P1# set routing-options srlg srlg-a srlg-cost 10
```

4. Configure MPLS on the interfaces and associate the SRLG with interface **ge-0/0/2.0** for the P1>PE2 link.

```
[edit protocols mpls]
user@P1# set interface ge-0/0/1.0
user@P1# set interface ge-0/0/2.0 srlg srlg-a
user@P1# set interface ge-0/0/3.0
user@P1# set interface ge-0/0/4.0
```

5. Enable RSVP on the interfaces and include the **link-protection** statement with the **exclude-srlg** option for interface **ge-0/0/2.0**.

```
[edit protocols rsvp]
user@P1# set interface ge-0/0/1.0
user@P1# set interface ge-0/0/2.0 link-protection exclude-srlg
user@P1# set interface ge-0/0/3.0
user@P1# set interface ge-0/0/4.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show protocols mpls**, **show protocols rsvp**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P1# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.12.2/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 192.168.27.2/24;
    }
    family mpls;
  }
}
ge-0/0/3 {
  unit 0 {
    family inet {
```

```
        address 192.168.23.2/24;
    }
    family mpls;
}
}
ge-0/0/4 {
    unit 0 {
        family inet {
            address 192.168.25.2/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.0.2/32;
        }
    }
}
```

```
user@P1# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
    interface ge-0/0/1.0;
    interface ge-0/0/2.0;
    interface ge-0/0/3.0;
    interface ge-0/0/4.0;
    interface lo0.0;
}
```

```
user@P1# show protocols mpls
interface ge-0/0/1.0;
interface ge-0/0/2.0 {
    srlg srlg-a;
}
interface ge-0/0/3.0;
interface ge-0/0/4.0;
```

```
user@P1# show protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0 {
    link-protection {
        exclude-srlg;
    }
}
interface ge-0/0/3.0;
interface ge-0/0/4.0;
}
```

```
user@P1# show routing-options
srlg {
    srlg-a {
        srlg-value 101;
        srlg-cost 10;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device P2

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure P2:

1. Configure the device interfaces.

```
[edit interfaces]
user@P2# set ge-0/0/1 unit 0 family inet address 192.168.13.3/24
user@P2# set ge-0/0/1 unit 0 family mpls
user@P2# set ge-0/0/2 unit 0 family inet address 192.168.37.3/24
user@P2# set ge-0/0/2 unit 0 family mpls
user@P2# set ge-0/0/3 unit 0 family inet address 192.168.23.3/24
user@P2# set ge-0/0/3 unit 0 family mpls
user@P2# set lo0 unit 0 family inet address 10.255.0.3/32
```

2. Configure OSPF on the interfaces.

```
[edit protocols ospf]
user@P2# set traffic-engineering
user@P2# set area 0.0.0.0 interface ge-0/0/1.0
user@P2# set area 0.0.0.0 interface ge-0/0/2.0
user@P2# set area 0.0.0.0 interface ge-0/0/3.0
user@P2# set area 0.0.0.0 interface lo0.0
```

3. Configure the SRLG definitions.

```
[edit routing-options]
user@P2# set routing-options srlg srlg-a srlg-value 101
user@P2# set routing-options srlg srlg-a srlg-cost 10
```

4. Configure MPLS on the interfaces and associate the SRLG with interface **ge-0/0/2.0** for the P2>PE2 link.

```
[edit protocols mpls]
user@P2# set interface ge-0/0/1.0
user@P2# set interface ge-0/0/2.0 srlg srlg-a
user@P2# set interface ge-0/0/3.0
```

5. Enable RSVP on the interfaces.

```
[edit protocols rsvp]
user@P2# set interface ge-0/0/1.0
user@P2# set interface ge-0/0/2.0
user@P2# set interface ge-0/0/3.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show protocols mpls**, **show protocols rsvp**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P2# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
```

```
        address 192.168.13.3/24;
    }
    family mpls;
}
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 192.168.37.3/24;
        }
        family mpls;
    }
}
ge-0/0/3 {
    unit 0 {
        family inet {
            address 192.168.23.3/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 10.255.0.3/32;
        }
    }
}
}
```

```
user@P2# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
    interface ge-0/0/1.0;
    interface ge-0/0/2.0;
    interface ge-0/0/3.0;
    interface lo0.0;
}
```

```
user@P2# show protocols mpls
interface ge-0/0/1.0;
interface ge-0/0/2.0 {
    srlg srlg-a;
}
interface ge-0/0/3.0;
}
```

```
user@P2# show protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;
```

```
user@P2# show routing-options
srlg {
    srlg-a {
        srlg-value 101;
        srlg-cost 10;
```

```
}
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device P3

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure P3:

1. Configure the device interfaces.

```
[edit interfaces]
user@P3# set ge-0/0/1 unit 0 family inet address 192.168.14.4/24
user@P3# set ge-0/0/1 unit 0 family mpls
user@P3# set ge-0/0/2 unit 0 family inet address 192.168.45.4/24
user@P3# set ge-0/0/2 unit 0 family mpls
user@P3# set lo0 unit 0 family inet address 10.255.0.4/32
```

2. Configure OSPF on the interfaces.

```
[edit protocols ospf]
user@P3# set traffic-engineering
user@P3# set area 0.0.0.0 interface ge-0/0/1.0
user@P3# set area 0.0.0.0 interface ge-0/0/2.0
user@P3# set area 0.0.0.0 interface lo0.0
```

3. Configure the SRLG definitions.

```
[edit routing-options]
user@P3# set routing-options srlg srlg-a srlg-value 101
user@P3# set routing-options srlg srlg-a srlg-cost 10
```

4. Configure MPLS on the interfaces.

```
[edit protocols mpls]
user@P3# set interface ge-0/0/1.0
user@P3# set interface ge-0/0/2.0
```

5. Enable RSVP on the interfaces.

```
[edit protocols rsvp]
user@P3# set interface ge-0/0/1.0
user@P3# set interface ge-0/0/2.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show protocols mpls**, **show protocols rsvp**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@P3# show interfaces
interfaces {
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.14.4/24;
```

```

    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 192.168.45.4/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.0.4/32;
    }
  }
}
}

user@P3# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface ge-0/0/1.0;
  interface ge-0/0/2.0;
  interface lo0.0;
}

user@P3# show protocols mpls
interface ge-0/0/1.0;
interface ge-0/0/2.0;

user@P3# show protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0;

user@P3# show routing-options
srlg {
  srlg-a {
    srlg-value 101;
    srlg-cost 10;
  }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device P4

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure P4:

1. Configure the device interfaces.

[edit interfaces]

```
user@P4# set ge-0/0/1 unit 0 family inet address 192.168.45.5/24
```

```

user@P4# set ge-0/0/1 unit 0 family mpls
user@P4# set ge-0/0/2 unit 0 family inet address 192.168.56.5/24
user@P4# set ge-0/0/2 unit 0 family mpls
user@P4# set ge-0/0/3 unit 0 family inet address 192.168.25.5/24
user@P4# set ge-0/0/3 unit 0 family mpls
user@P4# set lo0 unit 0 family inet address 10.255.0.5/32

```

2. Configure OSPF on the interfaces.

```

[edit protocols ospf]
user@P4# set traffic-engineering
user@P4# set area 0.0.0.0 interface ge-0/0/1.0
user@P4# set area 0.0.0.0 interface ge-0/0/2.0
user@P4# set area 0.0.0.0 interface ge-0/0/3.0
user@P4# set area 0.0.0.0 interface lo0.0

```

3. Configure the SRLG definitions.

```

[edit routing-options]
user@P4# set routing-options srlg srlg-a srlg-value 101
user@P4# set routing-options srlg srlg-a srlg-cost 10

```

4. Configure MPLS on the interfaces.

```

[edit protocols mpls]
user@P4# set interface ge-0/0/1.0
user@P4# set interface ge-0/0/2.0
user@P4# set interface ge-0/0/3.0

```

5. Enable RSVP on the interfaces.

```

[edit protocols rsvp]
user@P4# set interface ge-0/0/1.0
user@P4# set interface ge-0/0/2.0
user@P4# set interface ge-0/0/3.0

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show protocols mpls**, **show protocols rsvp**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@P4# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.45.5/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 192.168.56.5/24;
    }
    family mpls;
  }
}

```

```
ge-0/0/3 {
  unit 0 {
    family inet {
      address 192.168.25.5/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.0.5/32;
    }
  }
}

user@P4# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
  interface ge-0/0/1.0;
  interface ge-0/0/2.0;
  interface ge-0/0/3.0;
  interface lo0.0;
}

user@P4# show protocols mpls
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

user@P4# show protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0;
interface ge-0/0/3.0;

user@P4# show routing-options
srlg {
  srlg-a {
    srlg-value 101;
    srlg-cost 10;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device P5

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure P5:

1. Configure the device interfaces.

[edit interfaces]

```
user@P5# set ge-0/0/1 unit 0 family inet address 192.168.56.6/24
user@P5# set ge-0/0/1 unit 0 family mpls
user@P5# set ge-0/0/2 unit 0 family inet address 192.168.67.6/24
```



```

user@P5# set ge-0/0/2 unit 0 family mpls
user@P5# set lo0 unit 0 family inet address 10.255.0.6/32

```

2. Configure OSPF on the interfaces.

```

[edit protocols ospf]
user@P5# set traffic-engineering
user@P5# set area 0.0.0.0 interface ge-0/0/1.0
user@P5# set area 0.0.0.0 interface ge-0/0/2.0
user@P5# set area 0.0.0.0 interface lo0.0

```

3. Configure the SRLG definitions.

```

[edit routing-options]
user@P5# set routing-options srlg srlg-a srlg-value 101
user@P5# set routing-options srlg srlg-a srlg-cost 10

```

4. Configure MPLS on the interfaces.

```

[edit protocols mpls]
user@P5# set interface ge-0/0/1.0
user@P5# set interface ge-0/0/2.0

```

5. Enable RSVP on the interfaces.

```

[edit protocols rsvp]
user@P5# set interface ge-0/0/1.0
user@P5# set interface ge-0/0/2.0

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show protocols mpls**, **show protocols rsvp**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@P5# show interfaces
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.168.56.6/24;
    }
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 192.168.67.6/24;
    }
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.0.6/32;
    }
  }
}

```

```
user@P5# show protocols ospf
traffic-engineering;
area 0.0.0.0 {
    interface ge-0/0/1.0;
    interface ge-0/0/2.0;
    interface lo0.0;
}

user@P5# show protocols mpls
interface ge-0/0/1.0;
interface ge-0/0/2.0;

user@P5# show protocols rsvp
interface ge-0/0/1.0;
interface ge-0/0/2.0;

user@P5# show routing-options
srlg {
    srlg-a {
        srlg-value 101;
        srlg-cost 10;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Configuring Device PE2

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the *CLI User Guide*.

To configure PE2:

1. Configure the device interfaces.

[edit interfaces]

```
user@PE2# set ge-0/0/1 unit 0 family inet address 192.168.27.7/24
user@PE2# set ge-0/0/1 unit 0 family mpls
user@PE2# set ge-0/0/2 unit 0 family inet address 192.168.37.7/24
user@PE2# set ge-0/0/2 unit 0 family mpls
user@PE2# set ge-0/0/3 unit 0 family inet address 192.168.67.7/24
user@PE2# set ge-0/0/3 unit 0 family mpls
user@PE2# set lo0 unit 0 family inet address 10.255.0.7/32
```

2. Configure OSPF on the interfaces.

[edit protocols ospf]

```
user@PE2# set traffic-engineering
user@PE2# set area 0.0.0.0 interface ge-0/0/1.0
user@PE2# set area 0.0.0.0 interface ge-0/0/2.0
user@PE2# set area 0.0.0.0 interface ge-0/0/3.0
user@PE2# set area 0.0.0.0 interface lo0.0
```

3. Configure the SRLG definitions.

[edit routing-options]

```
user@PE2# set routing-options srlg srlg-a srlg-value 101
user@PE2# set routing-options srlg srlg-a srlg-cost 10
```

4. Configure MPLS on the interfaces.

```
[edit protocols mpls]
user@PE2# set interface ge-0/0/1.0
user@PE2# set interface ge-0/0/2.0
user@PE2# set interface ge-0/0/3.0
```

5. Enable RSVP on the interfaces.

```
[edit protocols rsvp]
user@PE2# set interface ge-0/0/1.0
user@PE2# set interface ge-0/0/2.0
user@PE2# set interface ge-0/0/3.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols ospf**, **show protocols mpls**, **show protocols rsvp**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE2# show interfaces
interfaces {
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 192.168.27.7/24;
      }
      family mpls;
    }
  }
  ge-0/0/2 {
    unit 0 {
      family inet {
        address 192.168.37.7/24;
      }
      family mpls;
    }
  }
  ge-0/0/3 {
    unit 0 {
      family inet {
        address 192.168.67.7/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.0.7/32;
      }
    }
  }
}

user@PE2# show protocols ospf
traffic-engineering;
```

```
area 0.0.0.0 {  
    interface ge-0/0/1.0;  
    interface ge-0/0/2.0;  
    interface ge-0/0/3.0;  
    interface lo0.0;  
}  
  
user@PE2# show protocols mpls  
interface ge-0/0/1.0;  
interface ge-0/0/2.0;  
interface ge-0/0/3.0;  
  
user@PE2# show protocols rsvp  
interface ge-0/0/1.0;  
interface ge-0/0/2.0;  
interface ge-0/0/3.0;  
  
user@PE2# show routing-options  
srlg {  
    srlg-a {  
        srlg-value 101;  
        srlg-cost 10;  
    }  
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying the SRLG Cost Is Added to the TE Link

Purpose Verify that the TE link is excluded if it belongs to the SRLG of the protected link when **link-protection** is configured with **exclude-srlg**. Issue the **show ted link detail** and **show rsvp session extensive bypass** commands on device P1.

Action user@P1> show ted link detail

```
...
10.255.0.2->192.168.27.7-1, Local: 192.168.27.2, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 1, StaticBW: 1000Mbps, AvailBW: 1000Mbps
  Color: 0 <none>
  SRLGs: srlg-a
    localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps
    localBW [4] 0bps [5] 0bps [6] 0bps [7] 0bps
[...]
```

user@P1> show rsvp session extensive bypass

```
Ingress RSVP: 1 sessions

10.255.0.7
  From: 10.255.0.2, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->192.168.27.7
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 299776
  Resv style: 1 SE, Label in: -, Label out: 299776
  Time left: -, Since: Fri Oct 21 13:19:21 2011
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 52081 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 1
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 192.168.25.5 (ge-0/0/4.0) 63 pkts
  RESV rcvfrom: 192.168.25.5 (ge-0/0/4.0) 63 pkts
  Explct route: 192.168.25.5 192.168.56.6 192.168.67.7
  Record route: <self> 192.168.25.5 192.168.56.6 192.168.67.7
Total 1 displayed, Up 1, Down 0
```

Meaning The shortest path for the bypass protecting the link P1>PE2 would have been P1>P2>PE2. Because the links P1>PE2 and P2>PE2 both belong to SRLG **srlg-a**, the link P2>PE2 is rejected for CSPF consideration due to the **exclude-srlg** constraint. Therefore, the computed path for the bypass becomes P1>P4>P5>PE2.

Related Documentation

- [SRLG Overview on page 80](#)
- [Example: Configuring SRLG on page 81](#)
- [Example: Configuring SRLG with Link Protection on page 95](#)
- [exclude-srlg on page 863](#)

Configuring the MPLS Transport Profile for OAM

- [MPLS Transport Profile Overview on page 136](#)
- [Example: Configuring the MPLS Transport Profile for OAM on page 136](#)

MPLS Transport Profile Overview

RFC 5654, *Requirements of an MPLS Transport Profile*, describes the requirements for the MPLS Transport Profile (MPLS-TP) that extends capabilities for Operation, Administration, and Maintenance (OAM) when MPLS is used for transport services and transport network operations. These capabilities help in troubleshooting and maintenance of a pseudowire or label-switched path (LSP).

MPLS-TP mechanisms for OAM contain two main components:

- Generic Associated Channel Label (GAL)—A special label that enables an exception mechanism that informs the egress label-switching router (LSR) that a packet it receives on an LSP belongs to an associated control channel or the control plane.
- Generic Associated Channel Header (G-Ach)—A special header field that identifies the type of payload contained in the MPLS label-switched paths (LSPs). G-Ach has the same format as a pseudowire associated control channel header.

For more information about MPLS-TP, see RFC 5654, *Requirements of an MPLS Transport Profile*. For specific information about GAL and G-Ach, see RFC 5586, *MPLS Generic Associated Channel*.

The following capabilities are supported in the Junos OS implementation of MPLS-TP:

- MPLS-TP OAM can send and receive packets with GAL and G-Ach, without IP encapsulation.
- Two unidirectional RSVP LSPs between a pair of routers can be associated with each other to create an associated bidirectional LSP for binding a path for the GAL and G-Ach OAM messages. A single Bidirectional Forwarding Detection (BFD) session is established for the associated bidirectional LSP.

Example: Configuring the MPLS Transport Profile for OAM

This example shows how to configure the MPLS Transport Profile (MPLS-TP) for sending and receiving of OAM GAL and G-Ach messages across a label-switched path (LSP).

- [Requirements on page 137](#)
- [Overview on page 137](#)
- [Configuration on page 139](#)
- [Verification on page 146](#)

Requirements

This example uses the following hardware and software components:

- Six devices that can be a combination of M Series, MX Series, and T Series routers
- Junos OS Release 12.1 or later running on the devices

Overview

Junos OS Release 12.1 and later support MPLS Transport Profile (MPLS-TP) Operation, Administration, and Maintenance (OAM) capabilities. MPLS-TP introduces new capabilities for OAM when MPLS is used for transport services and transport network operations. This includes configuring Generic Associated Channel Label (GAL) and Generic Associated Channel Header (G-Ach) for OAM messages.

This example shows how to configure MPLS-TP OAM capability to send and receive GAL and G-Ach OAM messages without IP encapsulation. In addition, it also shows how to associate two unidirectional RSVP label-switched paths (LSPs) between a pair of routers to create an associated bidirectional LSP for binding a path for the GAL and G-Ach OAM messages.

Junos OS Release 12.1 and later support the following MPLS-TP capabilities:

- MPLS-TP OAM capability and the infrastructure required for MPLS applications to send and receive packets with GAL and G-Ach, without IP encapsulation.
- LSP-ping and Bidirectional Forwarding Detection (BFD) applications to send and receive packets using GAL and G-Ach, without IP encapsulation on transport LSPs.
- The association of two unidirectional RSVP LSPs, between a pair of routers, with each other to create an associated bidirectional LSP for binding a path for the GAL and G-Ach OAM messages. The associated bidirectional LSP model is supported only for associating the primary paths. A single BFD session is established for the associated bidirectional LSP.

Junos OS Release 12.1 and later does not support the following MPLS-TP capabilities:

- Point-to-multipoint RSVP LSPs and BGP LSPs
- Loss Measurement and Delay Measurement

You can enable GAL and G-Ach OAM operation using the following configuration statements:

- **mpls-tp-mode**—Include this statement at the **[edit protocols mpls oam]** hierarchy level to enable GAL and G-Ach OAM operation, without IP encapsulation, on all LSPs in the MPLS network.

```
[edit protocols mpls oam]  
mpls-tp-mode;
```

Include this statement at the **[edit protocols mpls label-switched-path *lsp-name* oam]** hierarchy level to enable GAL and G-Ach OAM operation without IP encapsulation on a specific LSP in the network.

```
[edit protocols mpls label-switched-path lsp-name oam]
mpls-tp-mode;
```

- **associate-lsp *lsp-name* from *from-ip-address***—Include this statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level to configure associated bidirectional LSPs on the two ends of the LSP.

```
[edit protocols mpls label-switched-path lsp-name ]
associate-lsp lsp-name {
  from from-ip-address;
}
```

The **from *from-ip-address*** configuration for the LSP is optional. If omitted, it is derived from the **to** address of the ingress LSP configuration.

- **transit-lsp-association**—Include this statement at the **[edit protocols mpls]** hierarchy level to associate two LSPs at a transit router.


```
[edit protocols mpls]
transit-lsp-association transit-association-lsp-group-name {
  lsp-name-1 name-of-associated-lsp-1;
  from-1 address-of-associated-lsp-1;
  lsp-name-2 name-of-associated-lsp-2;
  from-2 address-of-associated-lsp-2;
}
```

The association of the LSPs in the transit nodes is useful for the return LSP path for TTL-expired LSP ping packets or traceroute.

In this example, R0 is the ingress router and R4 is the egress router. R1, R2, R3, and R5 are transit routers. The associated bidirectional LSP is established between the transit routers for sending and receiving the GAL and G-Ach OAM messages.

[Figure 21 on page 139](#) shows the topology used in this example.

Figure 21: MPLS-TP OAM Associated Bidirectional LSPs



g040928

Configuration

CLI Quick Configuration



NOTE: This example shows the configuration on all devices and shows step-by-step procedures for configuring the ingress router, R0, and transit router R1. Repeat the step-by-step procedure described for the ingress router, R0, on the egress router, R4. Repeat the step-by-step procedure for the transit router, R1, on the other transit routers, R2, R3, and R5. Be sure to modify the appropriate interface names, addresses, and other parameters appropriately.

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Router R0
set interfaces ge-4/1/1 unit 0 family inet address 10.10.11.1/30
set interfaces ge-4/1/1 unit 0 family iso
set interfaces ge-4/1/1 unit 0 family inet6
set interfaces ge-4/1/1 unit 0 family mpls
set interfaces ge-5/0/0 unit 0 family inet address 10.10.10.1/30
set interfaces ge-5/0/0 unit 0 family iso
set interfaces ge-5/0/0 unit 0 family inet6
set interfaces ge-5/0/0 unit 0 family mpls
set protocols rsvp interface ge-5/0/0.0
set protocols rsvp interface ge-4/1/1.0
set protocols mpls label-switched-path r0-to-r4 to 10.255.8.86
set protocols mpls label-switched-path r0-to-r4 oam mpls-tp-mode
set protocols mpls label-switched-path r0-to-r4 associate-lsp r4-to-r0 from 10.255.8.86
set protocols mpls interface ge-5/0/0.0
```

```
set protocols mpls interface ge-4/1/1.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-5/0/0.0
set protocols ospf area 0.0.0.0 interface ge-4/1/1.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

```
Router R1  set interfaces ge-0/0/5 unit 0 family inet address 10.10.10.2/30
set interfaces ge-0/0/5 unit 0 family iso
set interfaces ge-0/0/5 unit 0 family inet6
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces ge-0/2/2 unit 0 family inet address 10.10.12.2/30
set interfaces ge-0/2/2 unit 0 family iso
set interfaces ge-0/2/2 unit 0 family inet6
set interfaces ge-0/2/2 unit 0 family mpls
set interfaces ge-1/0/2 unit 0 family inet address 10.10.13.2/30
set interfaces ge-1/0/2 unit 0 family iso
set interfaces ge-1/0/2 unit 0 family inet6
set interfaces ge-1/0/2 unit 0 family mpls
set interfaces ge-2/0/2 unit 0 family inet address 10.10.11.2/30
set interfaces ge-2/0/2 unit 0 family iso
set interfaces ge-2/0/2 unit 0 family inet6
set interfaces ge-2/0/2 unit 0 family mpls
set protocols rsvp interface ge-0/2/2.0
set protocols rsvp interface ge-0/0/5.0
set protocols rsvp interface ge-1/0/2.0
set protocols rsvp interface ge-2/0/2.0
set protocols mpls transit-lsp-association trace1 lsp-name-1 r0-to-r4
set protocols mpls transit-lsp-association trace1 from-1 10.255.8.207
set protocols mpls transit-lsp-association trace1 lsp-name-2 r4-to-r0
set protocols mpls transit-lsp-association trace1 from-2 10.255.8.86
set protocols mpls interface ge-0/0/5.0
set protocols mpls interface ge-2/0/2.0
set protocols mpls interface ge-1/0/2.0
set protocols mpls interface ge-0/2/2.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0
set protocols ospf area 0.0.0.0 interface ge-0/2/2.0 metric 100
set protocols ospf area 0.0.0.0 interface ge-1/0/2.0
set protocols ospf area 0.0.0.0 interface ge-2/0/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

```
Router R2  set interfaces ge-0/2/3 unit 0 family inet address 10.10.13.1/30
set interfaces ge-0/2/3 unit 0 family iso
set interfaces ge-0/2/3 unit 0 family inet6
set interfaces ge-0/2/3 unit 0 family mpls
set interfaces ge-1/3/2 unit 0 family inet address 10.10.14.1/30
set interfaces ge-1/3/2 unit 0 family iso
set interfaces ge-1/3/2 unit 0 family inet6
set interfaces ge-1/3/2 unit 0 family mpls
set interfaces ge-1/3/4 unit 0 family inet address 10.10.15.1/30
set interfaces ge-1/3/4 unit 0 family iso
set interfaces ge-1/3/4 unit 0 family inet6
set interfaces ge-1/3/4 unit 0 family mpls
set protocols rsvp interface ge-0/2/3.0
set protocols rsvp interface ge-1/3/2.0
```

```

set protocols rsvp interface ge-1/3/4.0
set protocols mpls transit-lsp-association trace1 lsp-name-1 r0-to-r4
set protocols mpls transit-lsp-association trace1 from-1 10.255.8.207
set protocols mpls transit-lsp-association trace1 lsp-name-2 r4-to-r0
set protocols mpls transit-lsp-association trace1 from-2 10.255.8.86
set protocols mpls interface ge-0/2/3.0
set protocols mpls interface ge-1/3/2.0
set protocols mpls interface ge-1/3/4.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/2/3.0
set protocols ospf area 0.0.0.0 interface ge-1/3/2.0
set protocols ospf area 0.0.0.0 interface ge-1/3/4.0 metric 100
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

Router R3

```

set interfaces ge-1/2/1 unit 0 family inet address 10.10.16.2/30
set interfaces ge-1/2/1 unit 0 family iso
set interfaces ge-1/2/1 unit 0 family inet6
set interfaces ge-1/2/1 unit 0 family mpls
set interfaces ge-2/0/7 unit 0 family inet address 10.10.17.2/30
set interfaces ge-2/0/7 unit 0 family iso
set interfaces ge-2/0/7 unit 0 family inet6
set interfaces ge-2/0/7 unit 0 family mpls
set interfaces ge-2/2/0 unit 0 family inet address 10.10.14.2/30
set interfaces ge-2/2/0 unit 0 family iso
set interfaces ge-2/2/0 unit 0 family inet6
set interfaces ge-2/2/0 unit 0 family mpls
set protocols rsvp interface ge-2/2/0.0
set protocols rsvp interface ge-1/2/1.0
set protocols rsvp interface ge-2/0/7.0
set protocols mpls transit-lsp-association trace1 lsp-name-1 r0-to-r4
set protocols mpls transit-lsp-association trace1 from-1 10.255.8.207
set protocols mpls transit-lsp-association trace1 lsp-name-2 r4-to-r0
set protocols mpls transit-lsp-association trace1 from-2 10.255.8.86
set protocols mpls interface ge-2/2/0.0
set protocols mpls interface ge-1/2/1.0
set protocols mpls interface ge-2/0/7.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-2/2/0.0
set protocols ospf area 0.0.0.0 interface ge-1/2/1.0
set protocols ospf area 0.0.0.0 interface ge-2/0/7.0 metric 100
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

Router R4

```

set interfaces ge-0/0/3 unit 0 family inet address 10.10.16.1/30
set interfaces ge-0/0/3 unit 0 family iso
set interfaces ge-0/0/3 unit 0 family inet6
set interfaces ge-0/0/3 unit 0 family mpls
set protocols rsvp interface ge-0/0/3.0
set protocols mpls label-switched-path r4-to-r0 to 10.255.8.207
set protocols mpls label-switched-path r4-to-r0 oam mpls-tp-mode
set protocols mpls label-switched-path r4-to-r0 associate-lsp r0-to-r4 from 10.255.8.207
set protocols mpls interface ge-0/0/3.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/3.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

```

Router R5    set interfaces ge-1/2/0 unit 0 family inet address 10.10.15.2/30
              set interfaces ge-1/2/0 unit 0 family iso
              set interfaces ge-1/2/0 unit 0 family inet6
              set interfaces ge-1/2/0 unit 0 family mpls
              set interfaces ge-2/0/0 unit 0 family inet address 10.10.12.1/30
              set interfaces ge-2/0/0 unit 0 family iso
              set interfaces ge-2/0/0 unit 0 family inet6
              set interfaces ge-2/0/0 unit 0 family mpls
              set interfaces ge-4/0/7 unit 0 family inet address 10.10.17.1/30
              set interfaces ge-4/0/7 unit 0 family iso
              set interfaces ge-4/0/7 unit 0 family inet6
              set interfaces ge-4/0/7 unit 0 family mpls
              set protocols rsvp interface ge-2/0/0.0
              set protocols rsvp interface ge-1/2/0.0
              set protocols rsvp interface ge-4/0/7.0
              set protocols mpls transit-lsp-association trace1 lsp-name-1 r0-to-r4
              set protocols mpls transit-lsp-association trace1 from-1 10.255.8.207
              set protocols mpls transit-lsp-association trace1 lsp-name-2 r4-to-r0
              set protocols mpls transit-lsp-association trace1 from-2 10.255.8.86
              set protocols mpls interface ge-2/0/0.0
              set protocols mpls interface ge-1/2/0.0
              set protocols mpls interface ge-4/0/7.0
              set protocols ospf traffic-engineering
              set protocols ospf area 0.0.0.0 interface ge-2/0/0.0 metric 100
              set protocols ospf area 0.0.0.0 interface ge-1/2/0.0 metric 100
              set protocols ospf area 0.0.0.0 interface ge-4/0/7.0 metric 100
              set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

Configuring Device R0

Step-by-Step Procedure To configure the ingress router, R0:

1. Configure the interfaces.


```

[edit interfaces]
user@R0# set ge-4/1/1 unit 0 family inet address 10.10.11.1/30
user@R0# set ge-4/1/1 unit 0 family iso
user@R0# set ge-4/1/1 unit 0 family inet6
user@R0# set ge-4/1/1 unit 0 family mpls
user@R0# set ge-5/0/0 unit 0 family inet address 10.10.10.1/30
user@R0# set ge-5/0/0 unit 0 family iso
user@R0# set ge-5/0/0 unit 0 family inet6
user@R0# set ge-5/0/0 unit 0 family mpls

```
2. Configure MPLS on the interfaces.


```

[edit protocols mpls]
user@R0# set interface ge-5/0/0.0
user@R0# set interface ge-4/1/1.0

```
3. Configure an interior gateway protocol, such as OSPF.


```

[edit protocols ospf]
user@R0# set traffic-engineering
user@R0# set area 0.0.0.0 interface ge-5/0/0.0
user@R0# set area 0.0.0.0 interface ge-4/1/1.0
user@R0# set area 0.0.0.0 interface lo0.0 passive

```

4. Configure a signaling protocol, such as RSVP.

```
[edit protocols rsvp]
user@R0# set interface ge-5/0/0.0
user@R0# set interface ge-4/1/1.0
```

5. Configure the LSP.

```
[edit protocols mpls]
user@R0# set label-switched-path r0-to-r4 to 10.255.8.86
```

6. Enable GAL and G-Ach OAM operation without IP encapsulation on the LSPs.

```
[edit protocols mpls]
user@R0# set label-switched-path r0-to-r4 oam mpls-tp-mode
```

7. Configure associated bidirectional LSPs on the two ends of the LSP.

```
[edit protocols mpls]
user@R0# set label-switched-path r0-to-r4 associate-lsp to-r0 from 10.255.8.86
```

8. After you are done configuring the device, commit the configuration.

```
[edit]
user@R0# commit
```

Results Confirm your configuration by issuing the **show interfaces** and **show protocols** commands.

```
user@R0# show interfaces
ge-4/1/1 {
  unit 0 {
    family inet {
      address 10.10.11.1/30;
    }
    family iso;
    family inet6;
    family mpls;
  }
}
ge-5/0/0 {
  unit 0 {
    family inet {
      address 10.10.10.1/30;
    }
    family iso;
    family inet6;
    family mpls;
  }
}

user@R0# show protocols
rsvp {
  interface ge-5/0/0.0;
  interface ge-4/1/1.0;
}
mpls {
  label-switched-path r0-to-r4 {
    to 10.255.8.86;
    oam mpls-tp-mode;
```

```

        associate-lsp r4-to-r0 {
            from 10.255.8.86;
        }
    }
    interface ge-4/1/1.0;
    interface ge-5/0/0.0;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-5/0/0.0;
        interface ge-4/1/1.0;
        interface lo0.0 {
            passive;
        }
    }
}
}

```

Configuring Device R1

Step-by-Step Procedure To configure the transit router, R1:

1. Configure the interfaces.

[edit interfaces]

```

user@R1# set ge-0/0/5 unit 0 family inet address 10.10.10.2/30
user@R1# set ge-0/0/5 unit 0 family iso
user@R1# set ge-0/0/5 unit 0 family inet6
user@R1# set ge-0/0/5 unit 0 family mpls
user@R1# set ge-0/2/2 unit 0 family inet address 10.10.12.2/30
user@R1# set ge-0/2/2 unit 0 family iso
user@R1# set ge-0/2/2 unit 0 family inet6
user@R1# set ge-0/2/2 unit 0 family mpls
user@R1# set ge-2/0/2 unit 0 family inet address 10.10.11.2/30
user@R1# set ge-2/0/2 unit 0 family iso
user@R1# set ge-2/0/2 unit 0 family inet6
user@R1# set ge-2/0/2 unit 0 family mpls
user@R1# set ge-1/0/2 unit 0 family inet address 10.10.13.2/30
user@R1# set ge-1/0/2 unit 0 family iso
user@R1# set ge-1/0/2 unit 0 family inet6
user@R1# set ge-1/0/2 unit 0 family mpls

```

2. Configure MPLS on the interfaces.

[edit protocols mpls]

```

user@R1# set interface ge-0/0/5.0
user@R1# set interface ge-2/0/2.0
user@R1# set interface ge-1/0/2.0
user@R1# set interface ge-0/2/2.0

```

3. Configure an interior gateway protocol, such as OSPF.

[edit protocols ospf]

```

user@R1# set traffic-engineering
user@R1# set area 0.0.0.0 interface ge-0/0/5.0
user@R1# set area 0.0.0.0 interface ge-2/0/2.0
user@R1# set area 0.0.0.0 interface ge-1/0/2.0
user@R1# set area 0.0.0.0 interface ge-0/2/2.0 metric 100

```

```
user@R1# set area 0.0.0.0 interface lo0.0 passive
```

4. Configure a signaling protocol, such as RSVP.

```
[edit protocols rsvp]
user@R1# set interface ge-0/0/5.0
user@R1# set interface ge-2/0/2.0
user@R1# set interface ge-1/0/2.0
user@R1# set interface ge-0/2/2.0
```

5. Configure the association of the two LSPs on the transit router.

```
[edit protocols mpls]
user@R1# set transit-lsp-association trace1 lsp-name-1 r0-to-r4
user@R1# set transit-lsp-association trace1 from-1 10.255.8.207
user@R1# set transit-lsp-association trace1 lsp-name-2 r4-to-r0
user@R1# set transit-lsp-association trace1 from-2 10.255.8.86
```

6. If you are done configuring the device, commit the configuration.

```
[edit]
user@R1# commit
```

Results Confirm your configuration by issuing the **show interfaces** and **show protocols** commands.

```
user@R1# show interfaces
ge-0/0/5 {
  unit 0 {
    family inet {
      address 10.10.10.2/30;
    }
    family iso;
    family inet6;
    family mpls;
  }
}
ge-0/2/2 {
  unit 0 {
    family inet {
      address 10.10.12.2/30;
    }
    family iso;
    family inet6;
    family mpls;
  }
}
ge-2/0/2 {
  unit 0 {
    family inet {
      address 10.10.11.2/30;
    }
    family iso;
    family inet6;
    family mpls;
  }
}
ge-1/0/2 {
```

```
unit 0 {
  family inet {
    address 10.10.13.2/30;
  }
  family iso;
  family inet6;
  family mpls;
}

user@R1# show protocols
rsvp {
  interface ge-0/0/5.0;
  interface ge-2/0/2.0;
  interface ge-1/0/2.0;
  interface ge-0/2/2.0;
}
mpls {
  transit-lsp-association trace1 {
    lsp-name-1 r0-to-r4;
    from-1 10.255.8.207;
    lsp-name-2 r4-to-r0;
    from-2 10.255.8.86;
  }
  interface ge-0/0/5.0;
  interface ge-2/0/2.0;
  interface ge-1/0/2.0;
  interface ge-0/2/2.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-0/0/5.0;
    interface ge-1/0/2.0;
    interface ge-2/0/2.0;
    interface ge-0/2/2.0 {
      metric 100;
    }
    interface lo0.0 {
      passive;
    }
  }
}
}
```

Verification

Confirm that the configuration is working properly.

Verifying Associated Bidirectional LSPs

Purpose Verify that the associated bidirectional LSP configuration is working properly.

Action user@host> show mpls lsp

Ingress LSP: 1 sessions

To	From	State	Rt	P	ActivePath	LSPName
10.10.11.1	10.255.8.86	Up	0	*		r0-to-r4 Assoc-Bidir

Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

To	From	State	Rt	Style	Labelin	Labelout	LSPName
10.10.16.1	10.255.8.207	Up	0	1 FF	3		r4-to-r0 Assoc-Bidir

Total 2 displayed, Up 2, Down 0

Transit LSP: 1 sessions

To	From	State	Rt	Style	Labelin	Labelout	LSPName
10.10.10.2	10.255.8.168	Up	1	1 FF	301264	3	r0-to-r4 Assoc-Bidir

Total 3 displayed, Up 3, Down 0

user@host> show mpls lsp detail

Ingress LSP: 1 sessions

10.10.11.1

From: 10.255.8.86, State: Up, ActiveRoute: 0, LSPName: r0-to-r4

Associated Bidirectional

Associated LSP: r0-to-r4, 10.255.8.86

ActivePath: (primary)

LSPTYPE: Static Configured

LoadBalance: Random

Encoding type: Packet, Switching type: PSC-1, GPID: Unknown

*Primary State: Up

Egress LSP: 1 sessions

10.255.102.29

From: 10.255.102.172, LSPstate: Up, ActiveRoute: 0

LSPName: r4-to-r0, LSPpath: Primary

Associated Bidirectional

Associated LSP: 10.10.16.1, to-r0>

Suggested label received: -, Suggested label sent: -

Recovery label received: -, Recovery label sent: -

Resv style: 1 FF, Label in: 3, Label out: -

Time left: 144, Since: Fri Jun 17 21:41:05 2011

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500

Port number: sender 6 receiver 14468 protocol 0

PATH rcvfrom: 10.10.13.1 (ge-2/0/0.0) 84 pkts

Adspec: received MTU 1500

PATH sentto: localclient

RESV rcvfrom: localclient

Record route: 10.10.14.2 10.10.13.1 <self>

Transit LSP: 1 sessions

10.255.102.30

From: 10.255.102.172, LSPstate: Up, ActiveRoute: 1

LSPName: to_airstream, LSPpath: Primary

Associated Bidirectional

Associated LSP: r0-to-r4, 10.255.8.168

Suggested label received: -, Suggested label

Recovery label received: -, Recovery label sent: 3

Resv style: 1 FF, Label in: 301264, Label out: 3

Time left: 132, Since: Fri Jun 17 21:40:56 2011

```

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 28 receiver 14465 protocol 0
PATH rcvfrom: 10.10.13.1 (ge-2/0/0.0) 84 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.10.10.1 (ge-3/0/0.0) 84 pkts
RESV rcvfrom: 10.10.10.1 (ge-3/0/0.0) 84 pkts
Explot route: 10.10.10.1
Record route: 10.10.16.1 10.10.15.2 10.10.13.1 <self> 10.10.10.1

```

```

user@host> show mpls lsp bidirectional
Ingress LSP: 1 session
To          From          State Rt P    ActivePath    LSPname
10.255.8.86 10.255.8.207 Up    0 *          r0-to-r4
Assoc-Bidir
Total 1 displayed, Up 1, Down 0
Aug 28 06:56:26 [TRACE] [R0 coleman re0]
Egress LSP: 1 session
To          From          State Rt Style Labelin Labelout LSPname
10.255.8.207 10.255.8.86 Up    0 1 FF      3      - to-r0
Assoc-Bidir
Total 1 displayed, Up 1, Down 0
Aug 28 06:56:26 [TRACE] [R0 coleman re0]
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning The output of the `show mpls lsp`, `show mpls detail`, and `show mpls bidirectional` commands displays the details of the associated bidirectional LSPs and the LSP association information.

Configuring MPLS Inter-AS Link-Node Protection with Labeled BGP

- [Understanding MPLS Inter-AS Link Protection on page 148](#)
- [Example: Configuring MPLS Inter-AS Link-Node Protection on page 150](#)

Understanding MPLS Inter-AS Link Protection

Link protection is essential in an MPLS network to ensure traffic restoration in case of an interface failure. The ingress router chooses an alternate link through another interface to send traffic to its destination.


In [Figure 22 on page 149](#), autonomous system border routers (ASBRs) run external BGP (EBGP) to ASBRs in another autonomous system (AS) to exchange labels for /32 IPv4 routes. Inside the ASs, internal BGP (IBGP) propagates the routes to provider edge (PE) devices. If the link from Device ASBR3 to Device ASBR1 goes down, until Device ASBR3 reinstalls the new next hop, all traffic going toward AS 64510 from AS 64511 through the ASBR3-ASBR1 link is dropped. A fast traffic restoration can be achieved if Device ASBR3 preprograms a backup path either through Device ASBR4 or through a direct path to Device ASBR2 if one exists (not shown in the diagram). This assumes that Device ASBR3 learns a loop-free MPLS path for routes that need to be protected either through IBGP or EBGP.

This solution does not handle a failure on Device ASBR3 for traffic going toward AS 64511 from AS 64510 through the ASBR3-ASBR1 link. This solution is limited to downstream inter-AS link-node protection with labeled BGP. This solution does not support service

restoration between provider (P) and ASBR routers when there is an ASBR failure. For example, this solution does not handle a failure on the P3-ASBR3 link.

This supported functionality is similar to BGP multipath, except only one next hop is used for active forwarding, and a second path is in protected mode.

Figure 22: MPLS Inter-AS Link-Node Protection Conceptual Topology



In an MPLS inter-AS environment, link protection can be enabled when **labeled-unicast** is used to send traffic between ASs. Hence, MPLS inter-AS link protection is configured on the link between two routers in different ASs.

To configure link protection on an interface, use the **protection** statement at the **[edit protocols bgp group group-name family inet labeled-unicast]** hierarchy level:

```
protocols {
  bgp {
    group test1 {
      type external;
      local-address 192.168.1.2;
      family inet {
        labeled-unicast {
          protection;
        }
      }
    }
  }
}
```



NOTE: MPLS inter-AS link protection is supported only with labeled-unicast and external peers in a master routing instance.

The link on which protection is configured is known as the protection path. A protection path is selected only after the best path selection and is not selected in the following cases:

- The best path is a non-BGP path.
- Multiple next hops are active, as in BGP multipath.

Example: Configuring MPLS Inter-AS Link-Node Protection

This example shows how to configure tail-end protection in an inter-AS deployment with Layer 3 VPNs.

- [Requirements on page 150](#)
- [Overview on page 150](#)
- [Configuration on page 151](#)
- [Verification on page 161](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In [Figure 23 on page 151](#), autonomous system border routers (ASBRs) run external BGP (EBGP) to ASBRs in another autonomous system (AS) to exchange labels for /32 IPv4 routes. Inside the ASs, internal BGP (IBGP) propagates the routes to provider edge (PE) devices.

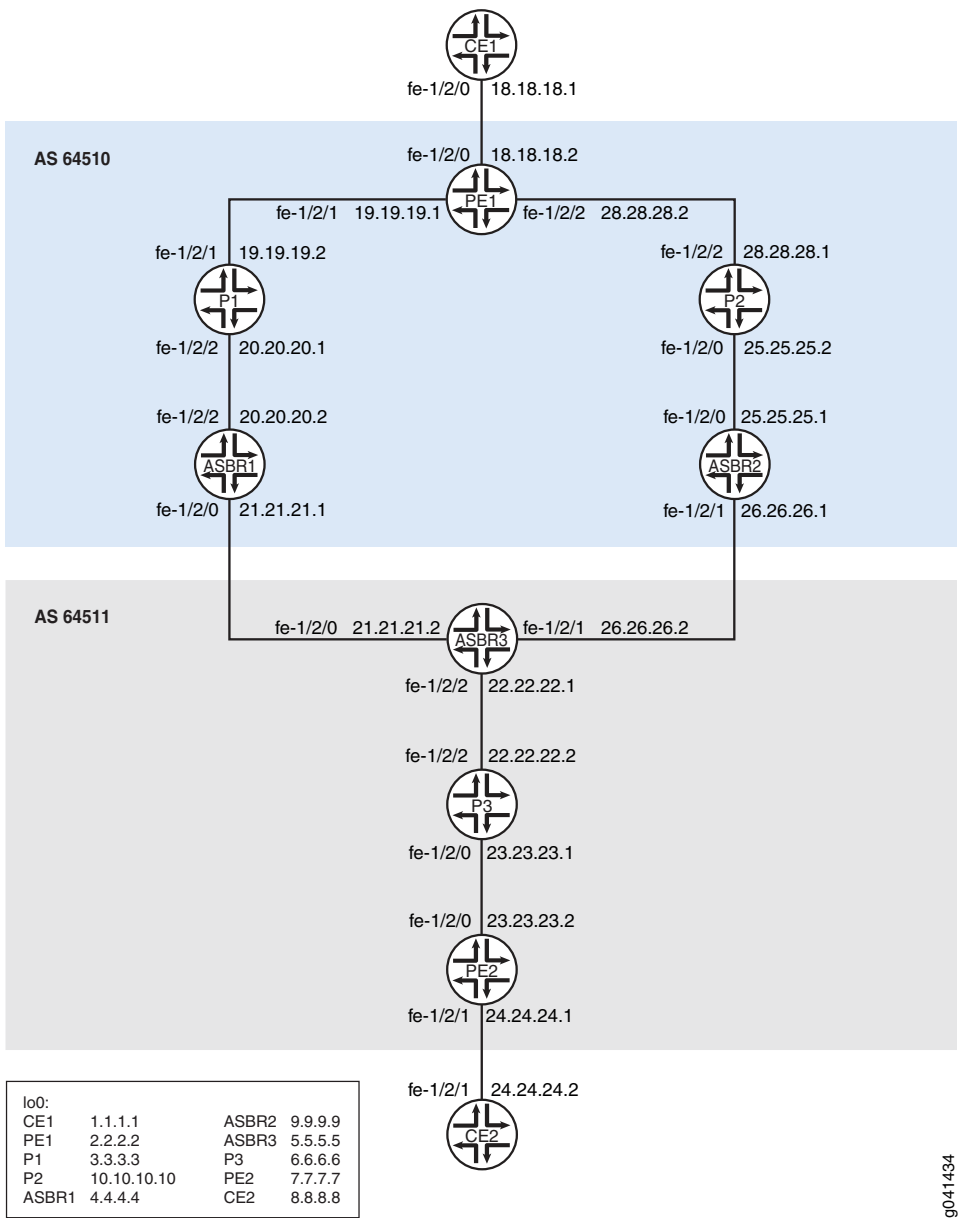
If the link from Device ASBR3 to Device ASBR1 goes down, until ASBR3 reinstalls the new next hop, all traffic going toward AS 64510 from AS 64511 through the ASBR3-ASBR1 link is dropped.

This example shows how to achieve fast traffic restoration by configuring Device ASBR3 to preprogram a backup path through Device ASBR2.



NOTE: This solution does not handle the Device P3 to Device ASBR3 failure. Nor does it handle a failure on Device ASBR3 for traffic going toward AS 64511 from AS 64510 through the ASBR3-ASBR1 link. This traffic is dropped.

Figure 23: MPLS Inter-AS Link-Node Protection Example Topology



“CLI Quick Configuration” on page 151 shows the configuration for all of the devices in Figure 23 on page 151.

The section “Step-by-Step Procedure” on page 156 describes the steps on Device ASBR3.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

Device ASBR1

```
set interfaces fe-1/2/0 unit 6 family inet address 20.20.20.2/30
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces fe-1/2/0 unit 0 family inet address 21.21.21.1/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 4.4.4.4/32
set protocols rsvp interface fe-1/2/2.0
set protocols rsvp interface lo0.0
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path To_PE1 to 2.2.2.2
set protocols mpls interface fe-1/2/2.0
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface lo0.0
set protocols bgp group To-PE1 type internal
set protocols bgp group To-PE1 local-address 4.4.4.4
set protocols bgp group To-PE1 family inet unicast
set protocols bgp group To-PE1 family inet labeled-unicast
set protocols bgp group To-PE1 export next-hop-self
set protocols bgp group To-PE1 neighbor 2.2.2.2 family inet labeled-unicast
set protocols bgp group To-ASBR3 type external
set protocols bgp group To-ASBR3 family inet labeled-unicast
set protocols bgp group To-ASBR3 export To-ASBR3
set protocols bgp group To-ASBR3 neighbor 21.21.21.2 peer-as 64511
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set policy-options policy-statement To-ASBR3 term 1 from route-filter 2.2.2.2/32 exact
set policy-options policy-statement To-ASBR3 term 1 then accept
set policy-options policy-statement To-ASBR3 term 2 then reject
set policy-options policy-statement next-hop-self then next-hop self
set routing-options autonomous-system 64510
```

Device ASBR2

```
set interfaces fe-1/2/0 unit 0 description to-P2
set interfaces fe-1/2/0 unit 0 family inet address 25.25.25.1/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 description to-ASBR3
set interfaces fe-1/2/1 unit 0 family inet address 26.26.26.1/30
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 9.9.9.9/32
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface lo0.0
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path To_PE1 to 2.2.2.2
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/1.0
set protocols mpls interface lo0.0
set protocols bgp group To-PE1 type internal
set protocols bgp group To-PE1 local-address 9.9.9.9
set protocols bgp group To-PE1 family inet unicast
set protocols bgp group To-PE1 family inet labeled-unicast
set protocols bgp group To-PE1 export next-hop-self
set protocols bgp group To-PE1 neighbor 2.2.2.2 family inet labeled-unicast
set protocols bgp group To-ASBR3 type external
set protocols bgp group To-ASBR3 family inet labeled-unicast
set protocols bgp group To-ASBR3 export To-ASBR3
set protocols bgp group To-ASBR3 neighbor 26.26.26.2 peer-as 64511
```

```

set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set policy-options policy-statement To-ASBR3 term 1 from route-filter 2.2.2.2/32 exact
set policy-options policy-statement To-ASBR3 term 1 then accept
set policy-options policy-statement To-ASBR3 term 2 then reject
set policy-options policy-statement next-hop-self then next-hop self
set routing-options autonomous-system 64510

```

Device ASBR3

```

set interfaces fe-1/2/0 unit 0 description to-ASBR1
set interfaces fe-1/2/0 unit 0 family inet address 21.21.21.2/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 description to-P3
set interfaces fe-1/2/2 unit 0 family inet address 22.22.22.1/30
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 description to-ASBR2
set interfaces fe-1/2/1 unit 0 family inet address 26.26.26.2/30
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 5.5.5.5/32
set protocols rsvp interface fe-1/2/2.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface fe-1/2/1.0
set protocols mpls traffic-engineering bgp-igp-both-ribs
set protocols mpls label-switched-path To_PE2 to 7.7.7.7
set protocols mpls interface lo0.0
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/2.0
set protocols mpls interface fe-1/2/1.0
set protocols bgp group To-PE2 type internal
set protocols bgp group To-PE2 local-address 5.5.5.5
set protocols bgp group To-PE2 family inet unicast
set protocols bgp group To-PE2 export next-hop-self
set protocols bgp group To-PE2 neighbor 7.7.7.7 family inet labeled-unicast
set protocols bgp group To-ASBR1 type external
set protocols bgp group To-ASBR1 family inet labeled-unicast protection
set protocols bgp group To-ASBR1 export To-ASBR1
set protocols bgp group To-ASBR1 neighbor 21.21.21.1 peer-as 64510
set protocols bgp group To-ASBR2 type external
set protocols bgp group To-ASBR2 family inet labeled-unicast protection
set protocols bgp group To-ASBR2 export To-ASBR2
set protocols bgp group To-ASBR2 neighbor 26.26.26.1 peer-as 64510
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/2.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/1.0
set policy-options policy-statement To-ASBR1 term 1 from route-filter 7.7.7.7/32 exact
set policy-options policy-statement To-ASBR1 term 1 then accept
set policy-options policy-statement To-ASBR1 term 2 then reject
set policy-options policy-statement To-ASBR2 term 1 from route-filter 7.7.7.7/32 exact
set policy-options policy-statement To-ASBR2 term 1 then accept
set policy-options policy-statement To-ASBR2 term 2 then reject
set policy-options policy-statement next-hop-self then next-hop self
set routing-options autonomous-system 64511

```

Device CE1	set interfaces fe-1/2/0 unit 0 family inet address 18.18.18.1/30 set interfaces lo0 unit 0 family inet address 1.1.1.1/32 set protocols ospf area 0.0.0.2 interface fe-1/2/0.0 set protocols ospf area 0.0.0.2 interface lo0.0 passive
Device CE2	set interfaces fe-1/2/1 unit 0 family inet address 24.24.24.2/30 set interfaces lo0 unit 0 family inet address 8.8.8.8/32 set protocols bgp group To_PE2 neighbor 24.24.24.1 export myroutes set protocols bgp group To_PE2 neighbor 24.24.24.1 peer-as 64511 set policy-options policy-statement myroutes from protocol direct set policy-options policy-statement myroutes then accept set routing-options autonomous-system 64509
Device P1	set interfaces fe-1/2/1 unit 0 family inet address 19.19.19.2/30 set interfaces fe-1/2/1 unit 0 family mpls set interfaces fe-1/2/2 unit 0 family inet address 20.20.20.1/30 set interfaces fe-1/2/2 unit 0 family mpls set interfaces lo0 unit 0 family inet address 3.3.3.3/32 set protocols rsvp interface fe-1/2/1.0 set protocols rsvp interface fe-1/2/2.0 set protocols rsvp interface lo0.0 set protocols mpls interface fe-1/2/1.0 set protocols mpls interface fe-1/2/2.0 set protocols mpls interface lo0.0 set protocols ospf traffic-engineering set protocols ospf area 0.0.0.0 interface fe-1/2/1.0 set protocols ospf area 0.0.0.0 interface fe-1/2/2.0 set protocols ospf area 0.0.0.0 interface lo0.0 passive
Device P2	set interfaces fe-1/2/0 unit 0 description to-ASBR2 set interfaces fe-1/2/0 unit 0 family inet address 25.25.25.2/30 set interfaces fe-1/2/0 unit 0 family mpls set interfaces fe-1/2/2 unit 0 description to-PE1 set interfaces fe-1/2/2 unit 0 family inet address 28.28.28.1/30 set interfaces fe-1/2/2 unit 0 family mpls set interfaces lo0 unit 0 family inet address 10.10.10.10/32 set protocols rsvp interface fe-1/2/0.0 set protocols rsvp interface fe-1/2/2.0 set protocols rsvp interface lo0.0 set protocols mpls interface fe-1/2/0.0 set protocols mpls interface fe-1/2/2.0 set protocols mpls interface lo0.0 set protocols ospf traffic-engineering set protocols ospf area 0.0.0.0 interface fe-1/2/0.0 set protocols ospf area 0.0.0.0 interface fe-1/2/2.0 set protocols ospf area 0.0.0.0 interface lo0.0 passive
Device P3	set interfaces fe-1/2/2 unit 0 family inet address 22.22.22.2/30 set interfaces fe-1/2/2 unit 0 family mpls set interfaces fe-1/2/0 unit 0 family inet address 23.23.23.1/30 set interfaces fe-1/2/0 unit 0 family mpls set interfaces lo0 unit 0 family inet address 6.6.6.6/32 set protocols rsvp interface fe-1/2/2.0 set protocols rsvp interface fe-1/2/0.0 set protocols rsvp interface lo0.0


```

set protocols mpls interface fe-1/2/2.0
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface lo0.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/2.0
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive

```

```

Device PE1
set interfaces fe-1/2/0 unit 0 family inet address 18.18.18.2/30
set interfaces fe-1/2/1 unit 0 family inet address 19.19.19.1/30
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces fe-1/2/2 unit 0 description to-P2
set interfaces fe-1/2/2 unit 0 family inet address 28.28.28.2/30
set interfaces lo0 unit 0 family inet address 2.2.2.2/32
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fe-1/2/2.0
set protocols mpls label-switched-path To_ASBR1 to 4.4.4.4
set protocols mpls label-switched-path To_ASBR2 to 9.9.9.9
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface lo0.0
set protocols mpls interface fe-1/2/2.0
set protocols bgp group To_ASBR1 type internal
set protocols bgp group To_ASBR1 local-address 2.2.2.2
set protocols bgp group To_ASBR1 family inet labeled-unicast
set protocols bgp group To_ASBR1 neighbor 4.4.4.4 family inet labeled-unicast resolve-vpn
set protocols bgp group To_PE2 type external
set protocols bgp group To_PE2 multihop ttl 20
set protocols bgp group To_PE2 local-address 2.2.2.2
set protocols bgp group To_PE2 family inet-vpn unicast
set protocols bgp group To_PE2 neighbor 7.7.7.7 peer-as 64511
set protocols bgp group To_ASBR2 type internal
set protocols bgp group To_ASBR2 local-address 2.2.2.2
set protocols bgp group To_ASBR2 family inet labeled-unicast
set protocols bgp group To_ASBR2 neighbor 9.9.9.9 family inet labeled-unicast resolve-vpn
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ospf area 0.0.0.0 interface fe-1/2/2.0
set policy-options policy-statement bgp-to-ospf term 1 from protocol bgp
set policy-options policy-statement bgp-to-ospf term 1 then accept
set policy-options policy-statement bgp-to-ospf term 2 then reject
set policy-options policy-statement vpnexport term 1 from protocol ospf
set policy-options policy-statement vpnexport term 1 then community add test_comm
set policy-options policy-statement vpnexport term 1 then accept
set policy-options policy-statement vpnexport term 2 then reject
set policy-options policy-statement vpnimport term 1 from protocol bgp
set policy-options policy-statement vpnimport term 1 from community test_comm
set policy-options policy-statement vpnimport term 1 then accept
set policy-options policy-statement vpnimport term 2 then reject
set policy-options community test_comm members target:1:64510
set routing-instances vpn2CE1 instance-type vrf
set routing-instances vpn2CE1 interface fe-1/2/0.0
set routing-instances vpn2CE1 route-distinguisher 1:64510
set routing-instances vpn2CE1 vrf-import vpnimport
set routing-instances vpn2CE1 vrf-export vpnexport

```

```

set routing-instances vpn2CE1 protocols ospf export bgp-to-ospf
set routing-instances vpn2CE1 protocols ospf area 0.0.0.2 interface fe-1/2/0.0
set routing-options autonomous-system 64510

```

Device PE2

```

set interfaces fe-1/2/0 unit 0 family inet address 23.23.23.2/30
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 24.24.24.1/30
set interfaces lo0 unit 0 family inet address 7.7.7.7/32
set protocols rsvp interface fe-1/2/0.0
set protocols rsvp interface lo0.0
set protocols mpls label-switched-path To-ASBR3 to 5.5.5.5
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface lo0.0
set protocols bgp group To_ASBR3 type internal
set protocols bgp group To_ASBR3 local-address 7.7.7.7
set protocols bgp group To_ASBR3 family inet labeled-unicast
set protocols bgp group To_ASBR3 neighbor 5.5.5.5 family inet labeled-unicast resolve-vpn
set protocols bgp group To_PE1 type external
set protocols bgp group To_PE1 multihop ttl 20
set protocols bgp group To_PE1 local-address 7.7.7.7
set protocols bgp group To_PE1 family inet-vpn unicast
set protocols bgp group To_PE1 neighbor 2.2.2.2 peer-as 64510
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface fe-1/2/0.0
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set policy-options policy-statement vpnexport term 1 from protocol bgp
set policy-options policy-statement vpnexport term 1 then community add test_comm
set policy-options policy-statement vpnexport term 1 then accept
set policy-options policy-statement vpnexport term 2 then reject
set policy-options policy-statement vpnimport term 1 from protocol bgp
set policy-options policy-statement vpnimport term 1 from community test_comm
set policy-options policy-statement vpnimport term 1 then accept
set policy-options policy-statement vpnimport term 2 then reject
set policy-options community test_comm members target:1:64510
set routing-instances vpn2CE2 instance-type vrf
set routing-instances vpn2CE2 interface fe-1/2/1.0
set routing-instances vpn2CE2 route-distinguisher 1:64510
set routing-instances vpn2CE2 vrf-import vpnimport
set routing-instances vpn2CE2 vrf-export vpnexport
set routing-instances vpn2CE2 protocols bgp group To_CE2 peer-as 64509
set routing-instances vpn2CE2 protocols bgp group To_CE2 neighbor 24.24.24.2
set routing-options autonomous-system 64511

```

**Step-by-Step
Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the EBGp scenario:

1. Configure the router interfaces.

```

[edit interfaces]
user@ASBR3# set fe-1/2/0 unit 0 description to-ASBR1
user@ASBR3# set fe-1/2/0 unit 0 family inet address 21.21.21.2/30
user@ASBR3# set fe-1/2/0 unit 0 family mpls

```

```

user@ASBR3# set fe-1/2/2 unit 0 description to-P3
user@ASBR3# set fe-1/2/2 unit 0 family inet address 22.22.22.1/30
user@ASBR3# set fe-1/2/2 unit 0 family mpls

```

```

user@ASBR3# set fe-1/2/1 unit 0 description to-ASBR2
user@ASBR3# set fe-1/2/1 unit 0 family inet address 26.26.26.2/30
user@ASBR3# set fe-1/2/1 unit 0 family mpls

```

```

user@ASBR3# set lo0 unit 0 family inet address 5.5.5.5/32

```

2. Configure an interior gateway protocol (IGP), such as OSPF or IS-IS.

```

[edit protocols ospf]
user@ASBR3# set traffic-engineering

```

```

[edit protocols ospf area 0.0.0.0]
user@ASBR3# set interface fe-1/2/2.0
user@ASBR3# set interface lo0.0 passive
user@ASBR3# set interface fe-1/2/1.0

```

3. Configure the autonomous system (AS) number.

```

[edit routing-options]
user@ASBR3# set autonomous-system 64511

```

4. Configure the routing policy.

```

[edit policy-options policy-statement To-ASBR1]
user@ASBR3# set term 1 from route-filter 7.7.7.7/32 exact
user@ASBR3# set term 1 then accept
user@ASBR3# set term 2 then reject

```

```

[edit policy-options policy-statement To-ASBR2]
user@ASBR3# set term 1 from route-filter 7.7.7.7/32 exact
user@ASBR3# set term 1 then accept
user@ASBR3# set term 2 then reject

```

```

[edit policy-options policy-statement next-hop-self]
user@ASBR3# set then next-hop self

```

5. Configure the EBGP sessions.

```

[edit protocols bgp group To-ASBR1]
user@ASBR3# set type external
user@ASBR3# set family inet labeled-unicast protection
user@ASBR3# set export To-ASBR1
user@ASBR3# set neighbor 21.21.21.1 peer-as 64510

```

```

[edit protocols bgp group To-ASBR2]
user@ASBR3# set type external
user@ASBR3# set family inet labeled-unicast protection
user@ASBR3# set export To-ASBR2
user@ASBR3# set neighbor 26.26.26.1 peer-as 64510

```

6. Configure the IBGP sessions.

```

[edit protocols bgp group To-PE2]

```

```
user@ASBR3# set type internal
user@ASBR3# set local-address 5.5.5.5
user@ASBR3# set family inet unicast
user@ASBR3# set export next-hop-self
user@ASBR3# set neighbor 7.7.7.7 family inet labeled-unicast
```

7. Configure MPLS.

```
[edit protocols mpls]
user@ASBR3# set traffic-engineering bgp-igp-both-ribs
user@ASBR3# set label-switched-path To_PE2 to 7.7.7.7
user@ASBR3# set interface lo0.0
user@ASBR3# set interface fe-1/2/0.0
user@ASBR3# set interface fe-1/2/2.0
user@ASBR3# set interface fe-1/2/1.0
```

8. Configure a signaling protocol.

```
[edit protocols rsvp]
user@ASBR3# set interface fe-1/2/2.0
user@ASBR3# set interface lo0.0
user@ASBR3# set interface fe-1/2/0.0
user@ASBR3# set interface fe-1/2/1.0
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options**, commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@ASBR3# show interfaces
fe-1/2/0 {
  unit 0 {
    description to-ASBR1;
    family inet {
      address 21.21.21.2/30;
    }
    family mpls;
  }
}
fe-1/2/1 {
  unit 0 {
    description to-ASBR2;
    family inet {
      address 26.26.26.2/30;
    }
    family mpls;
  }
}
fe-1/2/2 {
  unit 0 {
    description to-P3;
    family inet {
      address 22.22.22.1/30;
    }
    family mpls;
  }
}
```

```
}
lo0 {
  unit 0 {
    family inet {
      address 5.5.5.5/32;
    }
  }
}

user@ASBR3# show protocols
rsvp {
  interface fe-1/2/2.0;
  interface lo0.0;
  interface fe-1/2/0.0;
  interface fe-1/2/1.0;
}
mpls {
  traffic-engineering bgp-igp-both-ribs;
  label-switched-path To_PE2 {
    to 7.7.7.7;
  }
  interface lo0.0;
  interface fe-1/2/0.0;
  interface fe-1/2/2.0;
  interface fe-1/2/1.0;
}
bgp {
  group To-PE2 {
    type internal;
    local-address 5.5.5.5;
    family inet {
      unicast;
    }
    export next-hop-self;
    neighbor 7.7.7.7 {
      family inet {
        labeled-unicast;
      }
    }
  }
  group To-ASBR1 {
    type external;
    family inet {
      labeled-unicast {
        protection;
      }
    }
    export To-ASBR1;
    neighbor 21.21.21.1 {
      peer-as 64510;
    }
  }
  group To-ASBR2 {
    type external;
    family inet {
      labeled-unicast {
```

```
        protection;
    }
}
export To-ASBR2;
neighbor 26.26.26.1 {
    peer-as 64510;
}
}
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface fe-1/2/2.0;
        interface lo0.0 {
            passive;
        }
        interface fe-1/2/1.0;
    }
}

user@ASBR3# show policy-options
policy-statement To-ASBR1 {
    term 1 {
        from {
            route-filter 7.7.7/32 exact;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}
policy-statement To-ASBR2 {
    term 1 {
        from {
            route-filter 7.7.7/32 exact;
        }
        then accept;
    }
    term 2 {
        then reject;
    }
}
policy-statement next-hop-self {
    then {
        next-hop self;
    }
}

user@ASBR3# show routing-options
autonomous-system 64511;
```

If you are done configuring the devices, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Checking the BGP Neighbor Sessions on page 161](#)
- [Checking the Routes on page 163](#)

Checking the BGP Neighbor Sessions

Purpose Verify that BGP protection is enabled.

```

Action user@ASBR3# show bgp neighbor 21.21.21.1
Peer: 21.21.21.1+58259 AS 64510 Local: 21.21.21.2+179 AS 64511
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Export: [ To-ASBR1 ]
  Options: <Preference AddressFamily PeerAS Refresh>
  Options: <Protection>
  Address families configured: inet-labeled-unicast
  Holdtime: 90 Preference: 170
NLRI configured with protection: inet-labeled-unicast
  Number of flaps: 0
  Peer ID: 4.4.4.4      Local ID: 5.5.5.5      Active Holdtime: 90
  Keepalive Interval: 30      Group index: 4      Peer index: 0
  BFD: disabled, down
  Local Interface: fe-1/2/0.0
  NLRI for restart configured on peer: inet-labeled-unicast
  NLRI advertised by peer: inet-labeled-unicast
  NLRI for this session: inet-labeled-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-labeled-unicast
  NLRI of received end-of-rib markers: inet-labeled-unicast
  NLRI of all end-of-rib markers sent: inet-labeled-unicast
  Peer supports 4 byte AS extension (peer-as 64510)
  Peer does not support Addpath
  Table inet.0 Bit: 10001
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:      2
    Received prefixes:    1
    Accepted prefixes:    1
    Suppressed due to damping: 0
    Advertised prefixes:  1
  Last traffic (seconds): Received 7    Sent 20    Checked 32
  Input messages: Total 170    Updates 2    Refreshes 0    Octets 3326
  Output messages: Total 167    Updates 1    Refreshes 0    Octets 3288
  Output Queue[0]: 0

user@ASBR3# show bgp neighbor 26.26.26.1
Peer: 26.26.26.1+61072 AS 64510 Local: 26.26.26.2+179 AS 64511
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Export: [ To-ASBR2 ]
  Options: <Preference AddressFamily PeerAS Refresh>
  Options: <Protection>
  Address families configured: inet-labeled-unicast
  Holdtime: 90 Preference: 170
NLRI configured with protection: inet-labeled-unicast
  Number of flaps: 0
  Peer ID: 9.9.9.9      Local ID: 5.5.5.5      Active Holdtime: 90
  Keepalive Interval: 30      Group index: 5      Peer index: 0
  BFD: disabled, down
  Local Interface: fe-1/2/1.0
  NLRI for restart configured on peer: inet-labeled-unicast
  NLRI advertised by peer: inet-labeled-unicast
  NLRI for this session: inet-labeled-unicast
  Peer supports Refresh capability (2)

```



```

Stale routes from peer are kept for: 300
Peer does not support Restarter functionality
NLRI that restart is negotiated for: inet-labeled-unicast
NLRI of received end-of-rib markers: inet-labeled-unicast
NLRI of all end-of-rib markers sent: inet-labeled-unicast
Peer supports 4 byte AS extension (peer-as 64510)
Peer does not support Addpath
Table inet.0 Bit: 10002
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      1
Last traffic (seconds): Received 21   Sent 9   Checked 42
Input messages: Total 170   Updates 2   Refreshes 0   Octets 3326
Output messages: Total 168   Updates 1   Refreshes 0   Octets 3307
Output Queue[0]: 0

```

Meaning The output shows that the **Protection** option is enabled for the EBGp peers, Device ASBR1 and Device ASBR2.

This is also shown with the **NLRI configured with protection: inet-labeled-unicast** screen output.

Checking the Routes

Purpose Make sure that the backup path is installed in the routing table.

```

Action user@ASBR3> show route 2.2.2.2
inet.0: 12 destinations, 14 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32          *[BGP/170] 01:36:25, MED 2, localpref 100
                   AS path: 64510 I, validation-state: unverified
                   > to 21.21.21.1 via fe-1/2/0.0, Push 299824
                   to 26.26.26.1 via fe-1/2/1.0, Push 299808
                   [BGP/170] 01:36:25, MED 2, localpref 100
                   AS path: 64510 I, validation-state: unverified
                   > to 26.26.26.1 via fe-1/2/1.0, Push 299808

```

Meaning The **show route** command displays active as well as backup paths to Device PE1.

Related Documentation

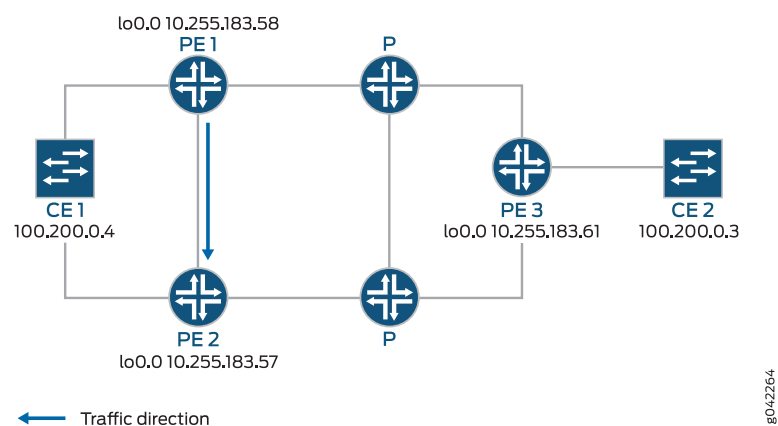
- *Example: Configuring Provider Edge Link Protection in Layer 3 VPNs*

Configuring Egress Protection Service Mirroring for BGP Signaled Layer 2 Services

If there is a link or node failure in the core network, a protection mechanism such as MPLS fast reroute can be triggered on the transport LSPs between the PE routers to repair the connection within tens of milliseconds. An egress protection LSP addresses the problem of a node-link failure at the edge of the network (for example, a failure of a PE router).

Figure 1 shows a simplified topology of the use case that explains this feature.

Figure 24: Egress Protection LSP Configured from Router PE1 to Router PE2



CE1 is multihomed to PE1 and PE2. There are two paths connecting CE1 and CE2. The working path is CE2-PE3-P-PE1-CE1, via pseudowire PW21. The protecting path is CE2-PE3-P-PE2-CE1, via pseudowire PW22. Traffic is flowing through the working path under normal circumstances. When the end-to-end OAM between CE1 and CE2 detects failure on the working path, traffic will be switched from the working path to the protecting path. The end-to-end failure detection and recovery relies on control plane hence should be relatively slow. To achieve faster protection, local repair mechanisms similar to those used by MPLS fast reroute should be used. In Figure 1 above, if link or node failed in the core network (like link failure on P-PE1, P-PE3, or node failure on P), the MPLS fast reroute will happen on the transport LSPs between PE1 and PE3. The failure could be locally repaired within tens of milliseconds. However, if link or node failure happens at the edge (like link failure on PE3-CE2 or node failure on PE3), there is no local repair currently so we have to rely on the CE1-CE2 end-to-end protection to repair the failure.

- Device CE2—Traffic origin
- Router PE3—Ingress PE router
- Router PE1— (Primary) Egress PE router
- Router PE2—Protector PE router
- Device CE1—Traffic destination

When the link between CE1– PE1 goes down, PE1 will briefly redirect that traffic towards CE1, to PE2. PE2 forwards it to CE1 until ingress router PE3 recalculates to forward the traffic to PE2.

Initially the traffic direction was; CE2 – PE3 – P – PE1 – CE1.

When the link between CE1– PE1 goes down, the traffic will be; CE2 – PE3 – P – PE1 – PE2 – CE1. PE3 then recalculates the path; CE2 – PE3 – P – PE2 – CE1.

1. Configure RSVP on PE1, PE2, and PE3.

```
[edit protocols]
user@PE1# set interface all
user@PE2# set interface all
user@PE3# set interface all
```

2. Configure MPLS.

```
[edit protocols mpls]
user@PE1# set interface all
user@PE2# set interface all
user@PE3# set interface all
```

3. Set PE1 as **primary** and PE2 as **protector** nodes.

```
[edit protocols mpls]
user@PE1# set egress-protection context-identifier address primary
user@PE2# set egress-protection context-identifier address protector
```

4. Enable **egress-protection** on PE1 and PE2.

```
[edit protocols bgp]
user@PE1# set group ibgp family l2vpn egress-protection
user@PE2# set group ibgp family l2vpn egress-protection
```

5. Configure LDP and ISIS on PE1, PE2, and PE3.

```
[edit protocols ldp]
user@PE1# set interface all
user@PE2# set interface all
user@PE3# set interface all

[edit protocols isis]
user@PE1# set interface all point-to-point
user@PE2# set interface all point-to-point
user@PE3# set interface all point-to-point
```

6. Configure a load balancing policy at PE1, PE2, and PE3.

```
[edit]
user@PE1# set policy-options policy-statement lb then load-balance per-packet
user@PE2# set policy-options policy-statement lb then load-balance per-packet
user@PE3# set policy-options policy-statement lb then load-balance per-packet
```

7. Configure the routing options at PE1, PE2, and PE3, to export routes based on the load balancing policy.

```
[edit]
user@PE1# set routing-options traceoptions file ro.log
user@PE1# set routing-options traceoptions flag normal
user@PE1# set routing-options traceoptions flag route
```

```

user@PE1# set routing-options autonomous-system 100
user@PE1# set routing-options forwarding-table export lb

```

```

[edit]
user@PE2# set routing-options traceoptions file ro.log
user@PE2# set routing-options traceoptions flag normal
user@PE2# set routing-options traceoptions flag route
user@PE2# set routing-options autonomous-system 100
user@PE2# set routing-options forwarding-table export lb

```

```

[edit]
user@PE3# set routing-options traceoptions file ro.log
user@PE3# set routing-options traceoptions flag normal
user@PE3# set routing-options traceoptions flag route
user@PE3# set routing-options autonomous-system 100
user@PE3# set routing-options forwarding-table export lb

```

8. Configure BGP at PE1 to advertise nrli from the routing instance with context-ID as next-hop.

```

[edit]
user@PE1# set routing-instances foo egress-protection context-identifier
context-identifier

```

9. Configure l2vpn at PE1, PE2, and PE3

At PE1:

```

[edit routing-instances]
foo {
  instance-type l2vpn;
  egress-protection {
    context-identifier {
      166.1.3.1;
    }
  }
  interface ge-2/0/2.0;
  route-distinguisher 10.255.183.58:1;
  vrf-target target:9000:1;
  protocols {
    l2vpn {
      encapsulation-type ethernet-vlan;
      site foo {
        site-identifier 1;
        multi-homing;
        site-preference primary;
        interface ge-2/0/2.0 {
          remote-site-id 2;
        }
      }
    }
  }
}

```

At PE2:

```

[edit routing-instances]
foo {
  instance-type l2vpn;

```

```

egress-protection {
  protector;
}
interface ge-2/0/2.0;
route-distinguisher 10.255.183.57:1;
vrf-target target:9000:1;
protocols {
  l2vpn {
    encapsulation-type ethernet-vlan;
    site foo {
      site-identifier 1;
      multi-homing;
      site-preference backup;
      interface ge-2/0/2.0 {
        remote-site-id 2;
      }
    }
  }
}
}

```

At PE3:

```

[edit routing-instances]
foo {
  instance-type l2vpn;
  interface ge-2/1/2.0;
  route-distinguisher 10.255.183.61:1;
  vrf-target target:9000:1;
  protocols {
    l2vpn {
      encapsulation-type ethernet-vlan;
      site foo {
        site-identifier 2;
        interface ge-2/1/2.0;
      }
    }
  }
}

```

Related Documentation

- [Configuring Per-Packet Load Balancing](#)
- [\[edit routing-instances\] Hierarchy Level](#)
- [Example: Configuring MPLS Egress Protection Service Mirroring for BGP Signaled Layer 2 Services on page 168](#)
- [Introduction to Configuring Layer 2 VPNs](#)
- [site-preference](#)

Example: Configuring MPLS Egress Protection Service Mirroring for BGP Signaled Layer 2 Services

If there is a link or node failure in the core network, a protection mechanism such as MPLS fast reroute can be triggered on the transport LSPs between the PE routers to repair the connection within tens of milliseconds. An egress protection LSP addresses the problem of a node-link failure at the edge of the network (for example, a failure of a PE router).

This example shows how to configure link protection for BGP signaled Layer 2 services.

- [Requirements on page 168](#)
- [Overview on page 168](#)
- [Configuration on page 169](#)
- [Verification on page 182](#)

Requirements

MX Series Routers running Junos OS Release 14.2 or later.

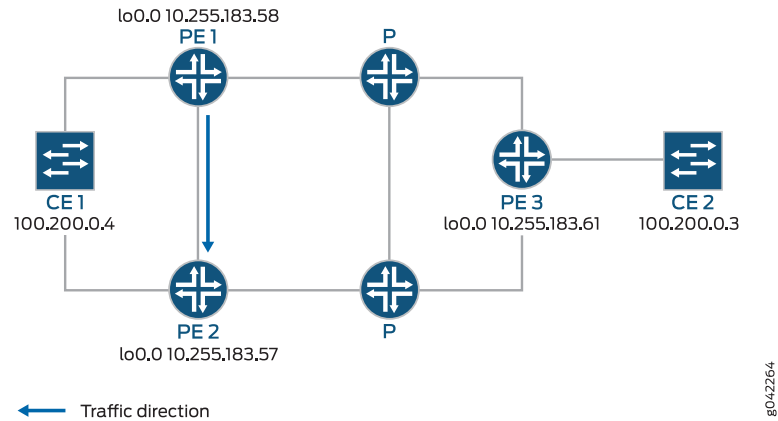
Overview

If there is a link or node failure in the core network, a protection mechanism such as MPLS fast reroute can be triggered on the transport LSPs between the PE routers to repair the connection within tens of milliseconds. An egress protection LSP addresses the problem of a node-link failure at the edge of the network (for example, a failure of a PE router).

This example includes the following configuration concepts and statements that are unique to the configuration of an egress protection LSP:

- **context-identifier**—Specifies an IPv4 or IPv6 address used to define the pair of PE routers participating in the egress protection LSP. It is assigned to each ordered pair of primary PE and the protector to facilitate protection establishment. This address is globally unique, or unique in the address space of the network where the primary PE and the protector reside.
- **egress-protection**—Configures the protector information for the protected Layer 2 circuit and configures the protector Layer 2 circuit at the **[edit protocols mpls]** hierarchy level. Configures an LSP as an egress protection LSP at the **[edit protocols mpls]** hierarchy level.
- **protector**—Configures the creation of standby pseudowires on the backup PE for link or node protection for the instance.

Figure 25: Egress Protection LSP Configured from Router PE1 to Router PE2



In the event of a failure of the egress PE Router PE1, traffic is switched to the egress protection LSP configured between Router PE1 and Router PE2 (the protector PE router):

- Device CE2—Traffic origin
- Router PE3—Ingress PE router
- Router PE1— (Primary) Egress PE router
- Router PE2—Protector PE router
- Device CE1—Traffic destination

When the link between CE1– PE1 goes down, PE1 will briefly redirect that traffic toward CE1, to PE2. PE2 forwards it to CE1 until ingress router PE3 recalculates to forward the traffic to PE2.

Initially the traffic direction was: CE2 – PE3 – P – PE1 – CE1.

When the link between CE1– PE1 goes down, the traffic will be: CE2 – PE3 – P – PE1 – PE2 –CE1. PE3 then recalculates the path: CE2 – PE3 – P – PE2 – CE1.

This example shows how to configure routers PE1, PE2, and PE3.

Configuration

- [Step-by-Step Procedure on page 172](#)
- [Results on page 176](#)

CLI Quick Configuration

To quickly configure an egress protection LSP, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configurations, copy and then paste the commands into the CLI and enter **commit** from configuration mode.

PE1

```
set protocols rsvp interface all
```

```
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls egress-protection context-identifier 166.1.3.1 primary
set protocols mpls egress-protection context-identifier 166.1.3.1 advertise-mode stub-alias
set protocols mpls egress-protection traceoptions file ep size 100m
set protocols mpls egress-protection traceoptions flag all
set protocols bgp traceoptions file bgp.log world-readable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.183.58
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family l2vpn signaling egress-protection
set protocols bgp group ibgp neighbor 100.200.0.3
set protocols bgp group ibgp neighbor 100.200.0.4
set protocols isis traceoptions file isis-edge size 10m world-readable
set protocols isis traceoptions flag error
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface all point-to-point
set protocols isis interface all level 2 metric 10
set protocols isis interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement lb then load-balance per-packet
set routing-options traceoptions file ro.log
set routing-options traceoptions flag all
set routing-options traceoptions flag route
set routing-options autonomous-system 100
set routing-options forwarding-table export lb
set routing-instances foo instance-type l2vpn
set routing-instances foo egress-protection context-identifier 166.1.3.1
set routing-instances foo interface ge-2/0/2.0
set routing-instances foo route-distinguisher 10.255.183.58:1
set routing-instances foo vrf-target target:9000:1
set routing-instances foo protocols l2vpn encapsulation-type ethernet-vlan
set routing-instances foo protocols l2vpn site foo site-identifier 1
set routing-instances foo protocols l2vpn site foo multi-homing
set routing-instances foo protocols l2vpn site foo site-preference primary
set routing-instances foo protocols l2vpn site foo interface ge-2/0/2.0 remote-site-id 2
```

PE2

```
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls egress-protection context-identifier 166.1.3.1 protector
set protocols mpls egress-protection context-identifier 166.1.3.1 advertise-mode stub-alias
set protocols mpls egress-protection traceoptions file ep size 100m
set protocols mpls egress-protection traceoptions flag all
set protocols bgp traceoptions file bgp.log world-readable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.183.57
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family l2vpn signaling egress-protection
set protocols bgp group ibgp neighbor 100.200.0.3
```



```

set protocols bgp group ibgp neighbor 100.200.0.4
set protocols isis traceoptions file isis-edge size 10m world-readable
set protocols isis traceoptions flag error
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface all point-to-point
set protocols isis interface all level 2 metric 10
set protocols isis interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement lb then load-balance per-packet
set routing-options traceoptions file ro.log
set routing-options traceoptions flag normal
set routing-options traceoptions flag route
set routing-options autonomous-system 100
set routing-options forwarding-table export lb
set routing-instances foo instance-type l2vpn
set routing-instances foo egress-protection protector
set routing-instances foo interface ge-2/0/2.0
set routing-instances foo route-distinguisher 10.255.183.57:1
set routing-instances foo vrf-target target:9000:1
set routing-instances foo protocols l2vpn encapsulation-type ethernet-vlan
set routing-instances foo protocols l2vpn site foo hot-standby
set routing-instances foo protocols l2vpn site foo site-identifier 1
set routing-instances foo protocols l2vpn site foo multi-homing
set routing-instances foo protocols l2vpn site foo site-preference backup
set routing-instances foo protocols l2vpn site foo interface ge-2/0/2.0 remote-site-id 2

```

PE3

```

set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp traceoptions file bgp.log world-readable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.183.61
set protocols bgp group ibgp family inet unicast
set protocols bgp group ibgp family l2vpn signaling
set protocols bgp group ibgp neighbor 100.200.0.3
set protocols bgp group ibgp neighbor 100.200.0.4
set protocols isis traceoptions file isis-edge size 10m world-readable
set protocols isis traceoptions flag error
set protocols isis level 1 disable
set protocols isis level 2 wide-metrics-only
set protocols isis interface all point-to-point
set protocols isis interface all level 2 metric 10
set protocols isis interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set policy-options policy-statement lb then load-balance per-packet
set routing-options traceoptions file ro.log
set routing-options traceoptions flag normal
set routing-options traceoptions flag route
set routing-options autonomous-system 100
set routing-options forwarding-table export lb

```

```

set routing-instances foo instance-type l2vpn
set routing-instances foo interface ge-2/1/2.0
set routing-instances foo route-distinguisher 10.255.183.61:1
set routing-instances foo vrf-target target:9000:1
set routing-instances foo protocols l2vpn encapsulation-type ethernet-vlan
set routing-instances foo protocols l2vpn site foo site-identifier 2
set routing-instances foo protocols l2vpn site foo interface ge-2/1/2.0 remote-site-id 1

```

Step-by-Step Procedure

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure an egress protection LSP for router PE1:

1. Configure RSVP.

```

[edit protocols rsvp]
user@PE1# set interface all
user@PE1# set interface fxp0.0 disable

```

2. Configure MPLS to use the egress protection LSP to protect against a link failure to Device CE1.

```

[edit protocols mpls]
user@PE1# set interface all
user@PE1# set interface fxp0.0 disable
user@PE1# set egress-protection context-identifier 166.1.3.1 primary
user@PE1# set egress-protection context-identifier 166.1.3.1 advertise-mode
stub-alias
user@PE1# set egress-protection traceoptions file ep size 100m
user@PE1# set egress-protection traceoptions flag all

```

3. Configure BGP.

```

[edit protocols bgp]
user@PE1# set traceoptions file bgp.log world-readable
user@PE1# set group ibgp type internal
user@PE1# set group ibgp local-address 10.255.183.58
user@PE1# set group ibgp family inet unicast
user@PE1# set group ibgp family l2vpn signaling egress-protection
user@PE1# set group ibgp neighbor 100.200.0.3
user@PE1# set group ibgp neighbor 100.200.0.4

```

4. Configure IS-IS.

```

[edit protocols isis]
user@PE1# set traceoptions file isis-edge size 10m world-readable
user@PE1# set traceoptions flag error
user@PE1# set level 1 disable
user@PE1# set level 2 wide-metrics-only
user@PE1# set interface all point-to-point
user@PE1# set interface all level 2 metric 10
user@PE1# set interface fxp0.0 disable

```

5. Configure LDP.

```

[edit protocols ldp]

```

- ```

user@PE1# set interface all
user@PE1# set interface fxp0.0 disable

```
6. Configure a load-balancing policy.

```

[edit]
user@PE1# set policy-options policy-statement lb then load-balance per-packet

```
  7. Configure the routing options to export routes based on the load-balancing policy.

```

[edit routing-options]
user@PE1# set traceoptions file ro.log
user@PE1# set traceoptions flag all
user@PE1# set autonomous-system 100
user@PE1# set forwarding-table export lb

```
  8. Configure BGP to advertise nrli from the routing instance with context-ID as next-hop.

```

[edit routing-instances]
user@PE1# set foo instance-type l2vpn
user@PE1# set foo egress-protection context-identifier 166.1.3.1
user@PE1# set foo interface ge-2/0/2.0
user@PE1# set foo route-distinguisher 10.255.183.58:1
user@PE1# set foo vrf-target target:9000:1

```
  9. Configure l2vpn instance to use the egress LSP configured.

```

[edit routing-instances]
user@PE1# set foo protocols l2vpn encapsulation-type ethernet-vlan
user@PE1# set foo protocols l2vpn site foo site-identifier 1
user@PE1# set foo protocols l2vpn site foo multi-homing
user@PE1# set foo protocols l2vpn site foo site-preference primary
user@PE1# set foo protocols l2vpn site foo interface ge-2/0/2.0 remote-site-id 2

```
  10. If you are done configuring the device, enter **commit** from configuration mode.

#### Step-by-Step Procedure

To configure an egress protection LSP for Router PE2:

1. Configure RSVP.

```

[edit protocols rsvp]
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable

```
2. Configure MPLS and the LSP that acts as the egress protection LSP.

```

[edit protocols mpls]
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable
user@PE2# set egress-protection context-identifier 166.1.3.1 protector
user@PE2# set egress-protection context-identifier 166.1.3.1 advertise-mode
stub-alias
user@PE2# set egress-protection traceoptions file ep size 100m
user@PE2# set egress-protection traceoptions flag all

```
3. Configure BGP.

```

[edit protocols bgp]
user@PE2# set traceoptions file bgp.log world-readable
user@PE2# set group ibgp type internal

```

- ```
user@PE2# set group ibgp local-address 10.255.183.57
user@PE2# set group ibgp family inet unicast
user@PE2# set group ibgp family l2vpn signaling
user@PE2# set group ibgp family l2vpn egress-protection
user@PE2# set group ibgp neighbor 100.200.0.3
user@PE2# set group ibgp neighbor 100.200.0.4
```
4. Configure IS-IS.

```
[edit protocols isis]
user@PE2# set traceoptions file isis-edge size 10m world-readable
user@PE2# set traceoptions flag error
user@PE2# set level 1 disable
user@PE2# set level 2 wide-metrics-only
user@PE2# set interface all point-to-point
user@PE2# set interface all level 2 metric 10
user@PE2# set interface fxp0.0 disable
```
 5. Configure LDP.

```
[edit protocols ldp]
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable
```
 6. Configure a load-balancing policy.

```
[edit]
user@PE2# set policy-options policy-statement lb then load-balance per-packet
```
 7. Configure the routing options to export routes based on the load-balancing policy.

```
[edit routing-options]
user@PE2# set traceoptions file ro.log
user@PE2# set traceoptions flag all
user@PE2# set autonomous-system 100
user@PE2# set forwarding-table export lb
```
 8. Configure BGP to advertise nrli from the routing instance with context-ID as next-hop.

```
[edit routing-instances]
user@PE2# set foo instance-type l2vpn
user@PE2# set foo egress-protection protector
user@PE2# set foo interface ge-2/0/2.0
user@PE2# set foo route-distinguisher 10.255.183.57:1
user@PE2# set foo vrf-target target:9000:1
```
 9. Configure l2vpn instance to use the egress LSP configured.

```
[edit routing-instances]
user@PE2# set foo protocols l2vpn encapsulation-type ethernet-vlan
user@PE2# set foo protocols l2vpn site foo hot-standby
user@PE2# set foo protocols l2vpn site foo site-identifier 1
user@PE2# set foo protocols l2vpn site foo multi-homing
user@PE2# set foo protocols l2vpn site foo site-preference backup
user@PE2# set foo protocols l2vpn site foo interface ge-2/0/2.0 remote-site-id 2
```
 10. If you are done configuring the device, enter **commit** from configuration mode.

**Step-by-Step
Procedure**

To configure an egress protection LSP for Router PE3:

1. Configure RSVP.

```
[edit protocols rsvp]
user@PE3# set interface all
user@PE3# set interface fxp0.0 disable
```
2. Configure MPLS.

```
[edit protocols mpls]
user@PE3# set interface all
user@PE3# set interface fxp0.0 disable
```
3. Configure BGP.

```
[edit protocols bgp]
user@PE3# set traceoptions file bgp.log world-readable
user@PE3# set group ibgp type internal
user@PE3# set group ibgp local-address 10.255.183.61
user@PE3# set group ibgp family inet unicast
user@PE3# set group ibgp family l2vpn signaling
user@PE3# set group ibgp neighbor 100.200.0.3
user@PE3# set group ibgp neighbor 100.200.0.4
```
4. Configure IS-IS.

```
[edit protocols isis]
user@PE3# set traceoptions file isis-edge size 10m world-readable
user@PE3# set traceoptions flag error
user@PE3# set level 1 disable
user@PE3# set level 2 wide-metrics-only
user@PE3# set protocols isis interface all point-to-point
[edit protocols isis]
user@PE3# set protocols isis interface all level 2 metric 10
[edit protocols isis]
user@PE3# set protocols isis interface fxp0.0 disable
```
5. Configure LDP.

```
[edit protocols ldp]
user@PE3# set interface all
user@PE3# set interface fxp0.0 disable
```
6. Configure a load-balancing policy.

```
[edit]
user@PE3# set policy-options policy-statement lb then load-balance per-packet
```
7. Configure the routing options to export routes based on the load-balancing policy.

```
[edit routing-options]
user@PE3# set traceoptions file ro.log
user@PE3# set traceoptions flag normal
user@PE3# set traceoptions flag route
user@PE3# set autonomous-system 100
user@PE3# set forwarding-table export lb
```
8. Configure BGP to advertise nlri from the routing instance with context-ID as next-hop.

```
[edit]
```

```
user@PE3# set routing-instances foo instance-type l2vpn
user@PE3# set routing-instances foo interface ge-2/1/2.0
user@PE3# set routing-instances foo route-distinguisher 10.255.183.61:1
user@PE3# set routing-instances foo vrf-target target:9000:1
```

9. Configure l2vpn to specify the interface that connects to the site and the remote interface to which you want the specified interface to connect.

```
[edit routing-instances]
user@PE3# set foo protocols l2vpn encapsulation-type ethernet-vlan
user@PE3# set foo protocols l2vpn site foo site-identifier 2
user@PE3# set foo protocols l2vpn site foo interface ge-2/1/2.0 remote-site-id 1
```

10. If you are done configuring the device, enter **commit** from configuration.

Results

From configuration mode, confirm your configuration on Router PE1 by entering the **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@PE1# show protocols
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
  egress-protection {
    context-identifier 166.1.3.1 {
      primary;
      advertise-mode stub-alias;
    }
    traceoptions {
      file ep size 100m;
      flag all;
    }
  }
}
bgp {
  traceoptions {
    file bgp.log world-readable;
  }
  group ibgp {
    type internal;
    local-address 10.255.183.58;
    family inet {
      unicast;
    }
  }
}
```

```

    family l2vpn {
        signaling {
            egress-protection;
        }
    }
    neighbor 100.200.0.3;
    neighbor 100.200.0.4;
}
isis {
    traceoptions {
        file isis-edge size 10m world-readable;
        flag error;
    }
    level 1 disable;
    level 2 wide-metrics-only;
    interface all {
        point-to-point;
        level 2 metric 10;
    }
    interface fxp0.0 {
        disable;
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}

```

```

[edit]
user@PE1# show policy-options
policy-statement lb {
    then {
        load-balance per-packet;
    }
}

```

```

[edit]
user@PE1# show routing-options
traceoptions {
    file ro.log;
    flag all;
}
autonomous-system 100;
forwarding-table {
    export lb;
}

```

```

[edit]
user@PE1# show routing-instances
foo {
    instance-type l2vpn;
    egress-protection {
        context-identifier {

```

```

        166.1.3.1;
    }
}
interface ge-2/0/2.0;
route-distinguisher 10.255.183.58:1;
vrf-target target:9000:1;
protocols {
  l2vpn {
    encapsulation-type ethernet-vlan;
    site foo {
      site-identifier 1;
      multi-homing;
      site-preference primary;
      interface ge-2/0/2.0 {
        remote-site-id 2;
      }
    }
  }
}
}
}

```

From configuration mode, confirm your configuration on Router PE2 by entering the **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

[edit]
user@PE2# show protocols
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
  egress-protection {
    context-identifier 166.1.3.1 {
      protector;
      advertise-mode stub-alias;
    }
    traceoptions {
      file ep size 100m;
      flag all;
    }
  }
}
}
bgp {
  traceoptions {
    file bgp.log world-readable;
  }
  group ibgp {
    type internal;
  }
}

```



```

    local-address 10.255.183.57;
    family inet {
        unicast;
    }
    family l2vpn {
        signaling {
            egress-protection;
        }
    }
    neighbor 100.200.0.3;
    neighbor 100.200.0.4;
}
}
isis {
    traceoptions {
        file isis-edge size 10m world-readable;
        flag error;
    }
    level 1 disable;
    level 2 wide-metrics-only;
    interface all {
        point-to-point;
        level 2 metric 10;
    }
    interface fxp0.0 {
        disable;
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}

[edit]
user@PE2# show policy-options
policy-statement lb {
    then {
        load-balance per-packet;
    }
}

[edit]
user@PE2# show routing-options
traceoptions {
    file ro.log;
    flag normal;
    flag route;
}
autonomous-system 100;
forwarding-table {
    export lb;
}

```

```
[edit]
user@PE2# show routing-instances
foo {
  instance-type l2vpn;
  egress-protection {
    protector;
  }
  interface ge-2/0/2.0;
  route-distinguisher 10.255.183.57:1;
  vrf-target target:9000:1;
  protocols {
    l2vpn {
      encapsulation-type ethernet-vlan;
      site foo {
        hot-standby;
        site-identifier 1;
        multi-homing;
        site-preference backup;
        interface ge-2/0/2.0 {
          remote-site-id 2;
        }
      }
    }
  }
}
```

From configuration mode, confirm your configuration on Router PE3 by entering the **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@PE3# show protocols
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
bgp {
  traceoptions {
    file bgp.log world-readable;
  }
  group ibgp {
    type internal;
    local-address 10.255.183.61;
    family inet {
      unicast;
    }
    family l2vpn {
```

```

        signaling;
    }
    neighbor 100.200.0.3;
    neighbor 100.200.0.4;
}
isis {
    traceoptions {
        file isis-edge size 10m world-readable;
        flag error;
    }
    level 1 disable;
    level 2 wide-metrics-only;
    interface all {
        point-to-point;
        level 2 metric 10;
    }
    interface fxp0.0 {
        disable;
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}

```

```

[edit]
user@PE3# show policy-options
policy-statement lb {
    then {
        load-balance per-packet;
    }
}

```

```

[edit]
user@PE3# show routing-options
traceoptions {
    file ro.log;
    flag normal;
    flag route;
}
autonomous-system 100;
forwarding-table {
    export lb;
}

```

```

[edit]
user@PE3# show routing-instances
foo {
    instance-type l2vpn;
    interface ge-2/1/2.0;
    route-distinguisher 10.255.183.61:1;
    vrf-target target:9000:1;
}

```

```

protocols {
  l2vpn {
    encapsulation-type ethernet-vlan;
    site foo {
      site-identifier 2;
      interface ge-2/1/2.0 {
        remote-site-id 1;
      }
    }
  }
}

```

Verification

Confirm that the configuration is working properly.

- [Verifying the L2VPN Configuration on page 182](#)
- [Verifying the Routing Instance Details on page 183](#)
- [Verifying the IS-IS Configuration on page 183](#)
- [Verifying the MPLS Configuration on page 184](#)

Verifying the L2VPN Configuration

Purpose Verify that LSP is protected by the connection protection logic.

Action From operational mode, run the **show l2vpn connections extensive** command.

```
user@PE2> show l2vpn connections extensive
```

```

Layer-2 VPN connections:
Legend for connection status (St)
EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch    WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down  NP -- interface hardware not present
CM -- control-word mismatch     -> -- only outbound connection is up
CN -- circuit not provisioned   <- -- only inbound connection is up
OR -- out of range             Up -- operational
OL -- no outgoing label        Dn -- down
LD -- local site signaled down  CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status IL -- no incoming label
MM -- MTU mismatch             MI -- Mesh-Group ID not available
BK -- Backup connection        ST -- Standby connection
PF -- Profile parse failure     PB -- Profile busy
RS -- remote site standby       SN -- Static Neighbor
LB -- Local site not best-site  RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status
Up -- operational
Dn -- down
Instance: foo
Local site: foo (1)
connection-site      Type  St  Time last up      # Up trans

```

```

2                               rmt  Up  Aug 3 00:08:14 2001          1
Local circuit: ge-2/0/2.0, Status: Up
Remote PE: 100.200.0.3
Incoming label: 32769, Outgoing label: 32768
Egress Protection: Yes
Time      Event      Interface/Lb1/PE
Aug 3 00:08:14 2001 PE route up
Aug 3 00:08:14 2001 Out 1b1 Update      32768
Aug 3 00:08:14 2001 In 1b1 Update      32769
Aug 3 00:08:14 2001 ckt0 up            fe-0/0/0.0

```

Meaning The **Egress Protection: Yes** output shows that the given PVC is protected by connection protection logic.

Verifying the Routing Instance Details

Purpose Verify the routing instance information and the context identifier configured on the primary, which is used as the next-hop address in case of node-link failure.

Action From operational mode, run the **show route foo detail** command.

```

user@PE2> show route foo detail

foo:
Router ID: 0.0.0.0
Type: l2vpn non-forwarding State: Active
Interfaces:
  lt-1/2/0.56
Route-distinguisher: 10.255.255.11:1
Vrf-import: [ __vrf-import-foo-internal__ ]
Vrf-export: [ __vrf-export-foo-internal__ ]
Vrf-import-target: [ target:100:200 ]
Vrf-export-target: [ target:100:200 ]
Fast-reroute-priority: low
Vrf-edge-protection-id: 166.1.3.1
Tables:
  foo.l2vpn.0      : 5 routes (3 active, 0 holddown, 0 hidden)
  foo.l2id.0       : 6 routes (2 active, 0 holddown, 0 hidden)

```

Meaning The context-id is set to **166.1.3.1** and the **Vrf-import: [__vrf-import-foo-internal__]** in the output mentions the policy used for rewriting the next-hop address.

Verifying the IS-IS Configuration

Purpose Verify the IS-IS context identifier information.

Action From operational mode, run the **show isis context-identifier detail** command.

```

user@PE2> show isis context-identifier detail

IS-IS context database:
Context      L  Owner      Role      Primary      Metric
166.1.3.1    2  MPLS      Protector  pro17-b-lr-R1 0
Advertiser pro17-b, Router ID 10.255.107.49, Level 2, tlv protector
Advertiser pro17-b-lr-R1, Router ID 10.255.255.11, Metric 1, Level 2, tlv prefix

```

Meaning Router PE2 is the protector and the configured context identifier is in use for the MPLS protocol.

Verifying the MPLS Configuration

Purpose Verify the context identifier details on the primary and protector PEs.

Action From operational mode, run the **show mpls context-identifier detail** command.

```
user@PE1> show mpls context-identifier detail
ID: 166.1.3.1
Type: primary, Metric: 1, Mode: alias
Total 1, Primary 1, Protector 0

user@PE2> show mpls context-identifier detail
ID: 166.1.3.1
Type: protector, Metric: 16777215, Mode: alias
Context table: __166.1.3.1__.mpls.0, Label out: 299968

user@PE2> show mpls egress-protection detail
```

```
Instance          Type      Protection-Type
foo               local-l2vpn Protector
Route Target 100:200
```

Meaning Context-id is **166.1.3.1**, advertise-mode is **alias**, the MPLS table created for egress protection is **__166.1.3.1__.mpls.0**, and the egress instance name is **foo**, which is of type **local-l2vpn**.

Related Documentation

- *Configuring Per-Packet Load Balancing*
- *[edit routing-instances] Hierarchy Level*
- *Introduction to Configuring Layer 2 VPNs*
- *site-preference*

Example: Configuring Layer 3 VPN Egress Protection with PLR as Protector

This example shows how to configure fast service restoration at the egress of a Layer 3 VPN when the customer is multihomed to the service provider.

Starting in Junos OS Release 15.1, enhanced PLR functionality is available, in which the PLR and the protector are co-located as one router. As part of this enhancement, there is no need to have a bypass LSP reroute traffic during local repair. Instead, the PLR or the protector can send the traffic directly to the target CE (in Co-located protector model where the PLR or the protector is also the backup PE that is directly connected to the

CE) or to the backup PE (in Centralized protector model where the backup PE is a separate router).

- [Requirements on page 185](#)
- [Overview on page 185](#)
- [Configuration on page 186](#)
- [Verification on page 202](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

This example requires Junos OS Release 15.1 or later.

Overview

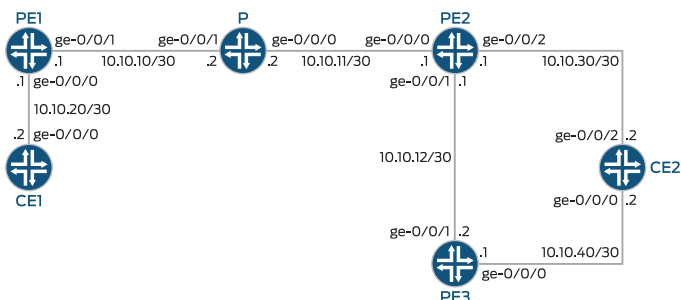
As a special scenario of egress node protection, if a router is both a Protector and a PLR, it installs backup next hops to protect the transport LSP. In particular, it does not need a bypass LSP for local repair.

In the Co-located protector model, the PLR or the Protector is directly connected to the CE via a backup AC, while in the Centralized protector model, the PLR or the protector has an MPLS tunnel to the backup PE. In either case, the PLR or the Protector will install a backup next hop with a label followed by a lookup in a **context label** table, i.e. **__context__.mpls.0**. When the egress node fails, the PLR or the Protector will switch traffic to this backup next hop in PFE. The outer label (the transport LSP label) of packets is popped, and the inner label (the layer 3 VPN label allocated by the egress node) is looked up in **__context__.mpls.0**, which results in forwarding the packets directly to the CE (in Collocated protector model) or the backup PE (in Centralized protector model).

Topology

Figure 26 on page 185 shows the sample network.

Figure 26: Co-located PLR and protector in collocated protector model



Configuration

- [Configuring Device CE1 on page 189](#)
- [Configuring Device PE1 on page 189](#)
- [Configuring Device P on page 191](#)
- [Configuring Device PE2 on page 192](#)
- [Configuring Device PE3 on page 193](#)
- [Configuring Device CE2 on page 195](#)
- [Results on page 195](#)

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device CE1	<pre> set interfaces ge-0/0/0 unit 0 family inet address 10.10.20.2/30 set interfaces lo0 unit 0 family inet address 10.255.162.87/32 </pre>
Device PE1	<pre> set interfaces ge-0/0/0 unit 0 family inet address 10.10.20.1/30 set interfaces ge-0/0/1 unit 0 family inet address 10.10.10.1/30 set interfaces ge-0/0/1 unit 0 family inet6 set interfaces ge-0/0/1 unit 0 family iso set interfaces ge-0/0/1 unit 0 family mpls set interfaces lo0 unit 0 family inet address 127.0.0.1/32 set interfaces lo0 unit 0 family inet address 10.255.162.84/32 primary set interfaces lo0 unit 0 family inet6 address abcd::10:255:162:84/128 primary set interfaces lo0 unit 0 family iso address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2084.00 set policy-options policy-statement vpn-exp term 1 from protocol direct set policy-options policy-statement vpn-exp term 1 from route filter 10.10.20.0/24 exact set policy-options policy-statement vpn-exp term 1 then community add vpn set policy-options policy-statement vpn-exp term 1 then accept set policy-options policy-statement vpn-imp term 1 from community vpn set policy-options policy-statement vpn-imp term 1 then accept set policy-options policy-statement vpn-imp term 2 then reject set policy-options community vpn members target:1:1 set routing-options autonomous-system 65000 set protocols rsvp interface all link-protection set protocols rsvp interface fxp0.0 disable set protocols mpls interface all set protocols mpls interface fxp0.0 disable set protocols bgp vpn-apply-export set protocols bgp group vpn type internal set protocols bgp group vpn local-address 10.255.162.84 set protocols bgp group vpn family inet-vpn unicast set protocols bgp group vpn neighbor 10.255.162.91 set protocols bgp group vpn neighbor 10.255.162.89 set protocols isis interface all set protocols isis interface fxp0.0 disable set protocols isis interface lo0.0 passive set routing-instances vpn instance-type vrf set routing-instances vpn interface ge-1/0/0.0 </pre>


```

set routing-instances vpn route-distinguisher 100:100
set routing-instances vpn vrf-import vpn-imp
set routing-instances vpn vrf-export vpn-exp
set routing-instances vpn vrf-table-label
set routing-instances vpn protocols bgp group vpn type external
set routing-instances vpn protocols bgp group vpn family inet unicast
set routing-instances vpn protocols bgp group vpn family inet6 unicast
set routing-instances vpn protocols bgp group vpn peer-as 65001
set routing-instances vpn protocols bgp group vpn as-override
set routing-instances vpn protocols bgp group vpn neighbor 10.10.20.2

```

Device P

```

set interfaces ge-0/0/0 unit 0 family inet address 10.10.11.2/30
set interfaces ge-0/0/0 unit 0 family inet6
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.10.10.2/30
set interfaces ge-0/0/1 unit 0 family inet6
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
set interfaces lo0 unit 0 family inet address 10.255.162.86/32 primary
set interfaces lo0 unit 0 family inet6 address abcd::10:255:162:86/128 primary
set interfaces lo0 unit 0 family iso address
    47.0005.80ff.f800.0000.0108.0001.0102.5516.2086.00
set protocols rsvp interface all link-protection
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis interface all
set protocols isis interface fxp0.0 disable

```

Device PE2

```

set interfaces ge-0/0/0 unit 0 family inet address 10.10.11.1/30
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family inet6
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 10.10.12.1/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family inet6
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 10.10.30.1/30
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
set interfaces lo0 unit 0 family inet address 10.255.162.91/32 primary
set interfaces lo0 unit 0 family iso address
    47.0005.80ff.f800.0000.0108.0001.0102.5516.2091.00
set interfaces lo0 unit 0 family inet6 address abcd::10:255:162:91/128 primary
set routing-options graceful-restart
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb
set protocols rsvp interface all link-protection
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path to_PE1 to 10.255.162.84
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls egress-protection context-identifier 1.1.1.1 protector
set protocols mpls egress-protection context-identifier 1.1.1.1 advertise-mode stub-alias

```

```

set protocols bgp vpn-apply-export
set protocols bgp group vpn type internal
set protocols bgp group vpn local-address 10.255.162.91
set protocols bgp group vpn family inet-vpn unicast egress-protection
set protocols bgp group vpn neighbor 10.255.162.84
set protocols bgp group vpn neighbor 10.255.162.89
set protocols isis traceoptions file isis.log
set protocols isis traceoptions flag all detail
set protocols isis level 2 disable
set protocols isis interface all
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 passive
set policy-options policy-statement pplb term 1 then load-balance per-packet
set policy-options policy-statement vpn-exp term 1 from protocol bgp
set policy-options policy-statement vpn-exp term 1 then community add vpn
set policy-options policy-statement vpn-exp term 1 then accept
set policy-options policy-statement vpn-imp term 1 from community vpn
set policy-options policy-statement vpn-imp term 1 then accept
set policy-options policy-statement vpn-imp term 2 then reject
set policy-options community vpn members target:1:1
set routing-instances vpn instance-type vrf
set routing-instances vpn interface ge-3/2/4.0
set routing-instances vpn route-distinguisher 100:100
set routing-instances vpn vrf-import vpn-imp
set routing-instances vpn vrf-export vpn-exp
set routing-instances vpn vrf-table-label
set routing-instances vpn protocols bgp group vpn type external
set routing-instances vpn protocols bgp group vpn family inet unicast
set routing-instances vpn protocols bgp group vpn family inet6 unicast
set routing-instances vpn protocols bgp group vpn peer-as 65001
set routing-instances vpn protocols bgp group vpn as-override
set routing-instances vpn protocols bgp group vpn neighbor 10.10.30.2

```

Device PE3

```

set interfaces ge-0/0/0 unit 0 family inet address 10.10.40.1/30
set interfaces ge-0/0/1 unit 0 family inet address 10.10.12.2/30
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family inet6
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
set interfaces lo0 unit 0 family inet address 10.255.162.89/32 primary
set interfaces lo0 unit 0 family iso address
    47.0005.80ff.f800.0000.0108.0001.0102.5516.2089.00
set interfaces lo0 unit 0 family inet6 address abcd::10:255:162:89/128 primary
set routing-options graceful-restart
set routing-options autonomous-system 65000
set routing-options forwarding-table export pplb
set protocols rsvp interface all link-protection
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path to_PE2 to 10.255.162.91
set protocols mpls label-switched-path to_PE1 to 10.255.162.84
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols mpls egress-protection context-identifier 1.1.1.1 primary
set protocols mpls egress-protection context-identifier 1.1.1.1 advertise-mode stub-alias
set protocols bgp vpn-apply-export
set protocols bgp group vpn type internal

```

```

set protocols bgp group vpn local-address 10.255.162.89
set protocols bgp group vpn family inet-vpn unicast
set protocols bgp group vpn neighbor 10.255.162.84 local-preference 300
set protocols bgp group vpn neighbor 10.255.162.91
set protocols isis level 2 disable
set protocols isis interface all
set protocols isis interface fxp0.0 disable
set protocols isis interface lo0.0 passive
set routing-instances vpn instance-type vrf
set routing-instances vpn egress-protection context-identifier 1.1.1.1
set routing-instances vpn interface ge-1/1/0.0
set routing-instances vpn route-distinguisher 100:100
set routing-instances vpn vrf-import vpn-imp
set routing-instances vpn vrf-export vpn-exp
set routing-instances vpn vrf-table-label
set routing-instances vpn protocols bgp group vpn type external
set routing-instances vpn protocols bgp group vpn family inet unicast
set routing-instances vpn protocols bgp group vpn family inet6 unicast
set routing-instances vpn protocols bgp group vpn peer-as 65001
set routing-instances vpn protocols bgp group vpn as-override
set routing-instances vpn protocols bgp group vpn neighbor 10.10.40.2

```

Device CE2

```

set interfaces ge-0/0/0 unit 0 family inet address 10.10.40.2/30
set interfaces ge-0/0/2 unit 0 family inet address 10.10.30.2/30
set interfaces lo0 unit 0 family inet address 127.0.0.1/32
set interfaces lo0 unit 0 family inet address 10.255.162.88/32 primary
set interfaces lo0 unit 0 family iso address
    47.0005.80ff.f800.0000.0108.0001.0102.5516.2088.00
set interfaces lo0 unit 0 family inet6 address abcd::10:255:162:88/128 primary

```

Configuring Device CE1

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure interfaces.


```

[edit interfaces]
user@CE1# set ge-0/0/0 unit 0 family inet address 10.10.20.2/30
user@CE1# set lo0 unit 0 family inet address 10.255.162.87/32

```

Configuring Device PE1

- Step-by-Step Procedure**
1. Configure the interfaces.


```

[edit interfaces]
user@PE1# set ge-0/0/0 unit 0 family inet address 10.10.20.1/30
user@PE1# set ge-0/0/1 unit 0 family inet address 10.10.10.1/30
user@PE1# set ge-0/0/1 unit 0 family iso
user@PE1# set ge-0/0/1 unit 0 family inet6
user@PE1# set ge-0/0/1 unit 0 family mpls
user@PE1# set lo0 unit 0 family inet address 127.0.0.1/32
user@PE1# set lo0 unit 0 family inet address 10.255.162.84/32 primary

```

- ```
user@PE1# set lo0 unit 0 family iso address
47.0005.80ff.f800.0000.0108.0001.0102.5516.2084.00
user@PE1# set lo0 unit 0 family inet6 address abcd::10:255:162:84/128 primary
```
2. Configure the autonomous system (AS) number.

```
[edit routing-options]
user@PE1# set autonomous-system 65000
user@PE1# set forwarding-table export pplb
```
  3. Configure RSVP.

```
[edit protocols rsvp]
user@PE1# set interface all link-protection
user@PE1# set interface fxp0.0 disable
```
  4. Enable MPLS.

```
[edit protocols mpls]
user@PE1# set interface all
user@PE1# set interface fxp0.0 disable
```
  5. Configure BGP.

```
[edit protocols bgp]
user@PE1# set group vpn type internal
user@PE1# set group vpn local-address 10.255.162.84
user@PE1# set group vpn family inet-vpn unicast
user@PE1# set group vpn neighbor 10.255.162.91
user@PE1# set group vpn neighbor 10.255.162.89
user@PE1# set vpn-apply-export
```
  6. Enable IS-IS.

```
[edit protocols isis]
user@PE1# set interface all
user@PE1# set interface fxp0.0 disable
user@PE1# set interface lo0.0 passive
```
  7. (Optional) Configure OSPF

```
[edit protocols ospf]
user@PE1# set area 0.0.0.0 interface all
user@PE1# set area 0.0.0.0 interface fxp0.0 disable
user@PE1# set area 0.0.0.0 interface lo0.0 passive
user@PE1# set traffic-engineering
```
  8. Configure the routing instance.

```
[edit routing-instances]
user@PE1# set vpn instance-type vrf
user@PE1# set vpn interface ge-1/0/0.0
user@PE1# set vpn route-distinguisher 100:100
user@PE1# set vpn vrf-import vpn-imp
user@PE1# set vpn vrf-export vpn-exp
user@PE1# set vpn vrf-table-label
user@PE1# set vpn protocols bgp group vpn type external
user@PE1# set vpn protocols bgp group vpn family inet unicast
user@PE1# set vpn protocols bgp group vpn family inet6 unicast
user@PE1# set vpn protocols bgp group vpn peer-as 65001
user@PE1# set vpn protocols bgp group vpn as-override
```

```
user@PE1# set vpn protocols bgp group vpn neighbor 10.10.20.2
```

9. Configure the routing policy.

```
[edit]
user@PE1# set policy-options policy-statement vpn-exp term 1 from protocol direct
user@PE1# set policy-options policy-statement vpn-exp term 1 from route filter
10.10.20.0/24 exact
user@PE1# set policy-options policy-statement vpn-exp term 1 then community
add vpn
user@PE1# set policy-options policy-statement vpn-exp term 1 then accept
user@PE1# set policy-options policy-statement vpn-imp term 1 from community
vpn
user@PE1# set policy-options policy-statement vpn-imp term 1 then accept
user@PE1# set policy-options policy-statement vpn-imp term 2 then reject
user@PE1# set policy-options community vpn members target:1:1
```

### Configuring Device P

#### Step-by-Step Procedure

1. Configure the device interfaces.

```
[edit interfaces]
user@P# set ge-0/0/0 unit 0 family inet address 10.10.11.2/30
user@P# set ge-0/0/0 unit 0 family inet6
user@P# set ge-0/0/0 unit 0 family iso
user@P# set ge-0/0/0 unit 0 family mpls
user@P# set ge-0/0/1 unit 0 family inet address 10.10.10.2/30
user@P# set ge-0/0/1 unit 0 family inet6
user@P# set ge-0/0/1 unit 0 family iso
user@P# set ge-0/0/1 unit 0 family mpls
user@P# set lo0 unit 0 family inet address 127.0.0.1/32
user@P# set lo0 unit 0 family inet address 10.255.162.86/32 primary
user@P# set lo0 unit 0 family inet6 address abcd::10:255:162:86/128 primary
user@P# set lo0 unit 0 family iso address
47.0005.80ff.f800.0000.0108.0001.0102.5516.2086.00
```

2. Enable IS-IS.

```
[edit protocols isis]
user@P# set interface all
user@P# set interface fxp0.0 disable
```

3. Enable MPLS.

```
[edit protocols mpls]
user@P# set interface all
user@P# set interface fxp0.0 disable
```

4. Configure RSVP.

```
[edit protocols rsvp]
user@P# set interface all link-protection
user@P# set interface fxp0.0 disable
```

5. (Optional) Configure OSPF.

```
[edit protocols ospf]
user@P# set area 0.0.0.0 interface all
user@P# set area 0.0.0.0 interface fxp0.0 disable
```

```

user@P# set area 0.0.0.0 interface lo0.0 passive
user@P# set traffic-engineering

```

## Configuring Device PE2

### Step-by-Step Procedure

1. Configure the interfaces.
 

```

[edit interfaces]
user@PE2# set ge-0/0/0 unit 0 family inet address 10.10.11.1/30
user@PE2# set ge-0/0/0 unit 0 family iso
user@PE2# set ge-0/0/0 unit 0 family inet6
user@PE2# set ge-0/0/0 unit 0 family mpls
user@PE2# set ge-0/0/1 unit 0 family inet address 10.10.12.1/30
user@PE2# set ge-0/0/1 unit 0 family iso
user@PE2# set ge-0/0/1 unit 0 family inet6
user@PE2# set ge-0/0/1 unit 0 family mpls
user@PE2# set ge-0/0/2 unit 0 family inet address 10.10.30.1/30
user@PE2# set lo0 unit 0 family inet address 127.0.0.1/32
user@PE2# set lo0 unit 0 family inet address 10.255.162.91/32 primary
user@PE2# set lo0 unit 0 family iso address
47.0005.80ff.f800.0000.0108.0001.0102.5516.2091.00
user@PE2# set lo0 unit 0 family inet6 address abcd::10:255:162:91/128 primary

```
2. Configure autonomous number(AS).
 

```

[edit routing-options]
user@PE2# set autonomous-system 65000
user@PE2# set forwarding-table export pplb

```
3. Configure RSVP.
 

```

[edit protocols rsvp]
user@PE2# set interface all link-protection
user@PE2# set interface fxp0.0 disable

```
4. Configure MPLS.
 

```

[edit protocols mpls]
user@PE2# set label-switched-path to_PE1 to 10.255.162.84
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable
user@PE2# set egress-protection context-identifier 1.1.1.1 protector
user@PE2# set egress-protection context-identifier 1.1.1.1 advertise-mode stub-alias

```
5. Configure BGP.
 

```

[edit protocols bgp]
user@PE2# set group vpn family inet-vpn unicast egress-protection
user@PE2# set group vpn local-address 10.255.162.91
user@PE2# set group vpn neighbor 10.255.162.84
user@PE2# set group vpn neighbor 10.255.162.89
user@PE2# set group vpn type internal
user@PE2# set vpn-apply-export

```
6. Configure IS-IS.
 

```

[edit protocols isis]
user@PE2# set interface all
user@PE2# set interface fxp0.0 disable

```

```

user@PE2# set interface lo0.0 passive
user@PE2# set level 2 disable
user@PE2# set traceoptions file isis.log
user@PE2# set traceoptions flag all detail

```

7. (Optional) Configure OSPF.

```

[edit protocols ospf]
user@PE2# set area 0.0.0.0 interface all
user@PE2# set area 0.0.0.0 interface fxp0.0 disable
user@PE2# set area 0.0.0.0 interface lo0.0 passive
user@PE2# set traffic-engineering

```

8. Configure the routing policy.

```

[edit policy-options]
user@PE2# set community vpn members target:1:1
user@PE2# set policy-statement pplb term 1 then load-balance per-packet
user@PE2# set policy-statement vpn-exp term 1 from protocol bgp
user@PE2# set policy-statement vpn-exp term 1 then community add vpn
user@PE2# set policy-statement vpn-exp term 1 then accept
user@PE2# set policy-statement vpn-imp term 1 from community vpn
user@PE2# set policy-statement vpn-imp term 1 then accept
user@PE2# set policy-statement vpn-imp term 2 then reject

```

9. Configure the routing instance.

```

[edit routing-instances]
user@PE2# set vpn instance-type vrf
user@PE2# set vpn interface ge-3/2/4.0
user@PE2# set vpn route-distinguisher 100:100
user@PE2# set vpn vrf-import vpn-imp
user@PE2# set vpn vrf-export vpn-exp
user@PE2# set vpn vrf-table-label
user@PE2# set vpn protocols bgp group vpn type external
user@PE2# set vpn protocols bgp group vpn family inet unicast
user@PE2# set vpn protocols bgp group vpn family inet6 unicast
user@PE2# set vpn protocols bgp group vpn peer-as 65001
user@PE2# set vpn protocols bgp group vpn as-override
user@PE2# set vpn protocols bgp group vpn neighbor 10.10.30.2

```

### Configuring Device PE3

#### Step-by-Step Procedure

1. Configure the interfaces.

```

[edit interfaces]
user@PE3# set ge-0/0/0 unit 0 family inet address 10.10.40.1/30
user@PE3# set ge-0/0/1 unit 0 family inet address 10.10.12.2/30
user@PE3# set ge-0/0/1 unit 0 family iso
user@PE3# set ge-0/0/1 unit 0 family inet6
user@PE3# set ge-0/0/1 unit 0 family mpls
user@PE3# set lo0 unit 0 family inet address 127.0.0.1/32
user@PE3# set lo0 unit 0 family inet address 10.255.162.89/32 primary
user@PE3# set lo0 unit 0 family iso address
47.0005.80ff.f800.0000.0108.0001.0102.5516.2089.00
user@PE3# set lo0 unit 0 family inet6 address abcd::10:255:162:89/128 primary

```

2. Configure the autonomous number (AS).

- ```
[edit routing-options]
user@PE3# set autonomous-system 65000
user@PE3# set forwarding-table export pplb
```
3. Configure RSVP.

```
[edit protocols rsvp]
user@PE3# set interface all link-protection
user@PE3# set interface fxp0.0 disable
```
 4. Configure MPLS.

```
[edit protocols mpls]
user@PE3# set interface all
user@PE3# set interface fxp0.0 disable
user@PE3# set egress-protection context-identifier 1.1.1.1 primary
user@PE3# set egress-protection context-identifier 1.1.1.1 advertise-mode stub-alias
user@PE3# set label-switched-path to_PE2 to 10.255.162.91
user@PE3# set label-switched-path to_PE1 to 10.255.162.84
```
 5. Configure BGP.

```
[edit protocols bgp]
user@PE3# set group vpn type internal
user@PE3# set group vpn local-address 10.255.162.89
user@PE3# set group vpn family inet-vpn unicast
user@PE3# set group vpn neighbor 10.255.162.84 local-preference 300
user@PE3# set group vpn neighbor 10.255.162.91
user@PE3# set vpn-apply-export
```
 6. Configure IS-IS.

```
[edit protocols isis]
user@PE3# set interface all
user@PE3# set interface fxp0.0 disable
user@PE3# set interface lo0.0 passive
user@PE3# set level 2 disable
```
 7. (Optional) Configure OSPF.

```
[edit protocols ospf]
user@PE3# set area 0.0.0.0 interface all
user@PE3# set area 0.0.0.0 interface fxp0.0 disable
user@PE3# set area 0.0.0.0 interface lo0.0 passive
user@PE3# set traffic-engineering
```
 8. Configure the routing instance.

```
[edit routing-instances]
user@PE3# set vpn egress-protection context-identifier 1.1.1.1
user@PE3# set vpn instance-type vrf
user@PE3# set vpn interface ge-1/1/0.0
user@PE3# set vpn protocols bgp group vpn type external
user@PE3# set vpn protocols bgp group vpn family inet unicast
user@PE3# set vpn protocols bgp group vpn family inet6 unicast
user@PE3# set vpn protocols bgp group vpn peer-as 65001
user@PE3# set vpn protocols bgp group vpn as-override
user@PE3# set vpn protocols bgp group vpn neighbor 10.10.40.2
user@PE3# set vpn route-distinguisher 100:100
user@PE3# set vpn vrf-export vpn-exp
```



```

user@PE3# set vpn vrf-import vpn-imp
user@PE3# set vpn vrf-table-label

```

Configuring Device CE2

Step-by-Step Procedure

1. Configure the interfaces.

```

[edit interfaces]
user@CE2# set ge-0/0/0 unit 0 family inet address 10.10.40.2/30
user@CE2# set ge-0/0/2 unit 0 family inet address 10.10.30.2/30
user@CE2# set lo0 unit 0 family inet address 127.0.0.1/32
user@CE2# set lo0 unit 0 family inet address 10.255.162.88/32 primary
user@CE2# set lo0 unit 0 family iso address
47.0005.80ff.f800.0000.0108.0001.0102.5516.2088.00
user@CE2# set lo0 unit 0 family inet6 address abcd::10:255:162:88/128 primary

```

Results

From configuration mode, confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

Device CE1 user@CE1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.20.2/30;
    }
  }
}

Device PE1 user@PE1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.20.1/30;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 10.10.10.1/30;
    }
    family iso;
    family inet6;
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 127.0.0.1/32;
      address 10.255.162.84/32 {
        primary;

```

```

    }
  }
  family iso {
    address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2084.00;
  }
  family inet6 {
    address abcd::10:255:162:84/128 {
      primary;
    }
  }
}

```

user@PE1# show protocols

```

rsvp {
  interface all {
    link-protection;
  }
  interface fxp0.0 {
    disable;
  }
}
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
bgp {
  vpn-apply-export;
  group vpn {
    type internal;
    local-address 10.255.162.84;
    family inet-vpn {
      unicast;
    }
    neighbor 10.255.162.91;
    neighbor 10.255.162.89;
  }
}
isis {
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface lo0.0 {
    passive;
  }
}

```

Device P

user@P# show interfaces

```

ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.11.2/30;
    }
  }
}

```

```

        family iso;
        family inet6;
        family mpls;
    }
}
ge-0/0/1 {
    unit 0 {
        family inet {
            address 10.10.10.2/30;
        }
        family iso;
        family inet6;
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 127.0.0.1/32;
            address 10.255.162.86/32 {
                primary;
            }
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2086.00;
        }
        family inet6 {
            address abcd::10:255:162:86/128 {
                primary;
            }
        }
    }
}

```

user@P# show protocols

```

rsvp {
    interface all {
        link-protection;
    }
    interface fxp0.0 {
        disable;
    }
}
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
isis {
    interface all;
    interface fxp0.0 {
        disable;
    }
}

```

```
Device PE2      user@PE2# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.11.1/30;
    }
    family iso;
    family inet6;
    family mpls;
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 10.10.12.1/30;
    }
    family iso;
    family inet6;
    family mpls;
  }
}
ge-0/0/2 {
  unit 0 {
    family inet {
      address 10.10.30.1/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 127.0.0.1/32;
      address 10.255.162.91/32 {
        primary;
      }
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2091.00;
    }
    family inet6 {
      address abcd::10:255:162:91/128 {
        primary;
      }
    }
  }
}

user@PE2# show protocols
rsvp {
  interface all {
    link-protection;
  }
  interface fxp0.0 {
    disable;
  }
}
```

```

}
mpls {
  label-switched-path to_PE1 {
    to 10.255.162.84;
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
  egress-protection {
    context-identifier 1.1.1.1 {
      protector;
      advertise-mode stub-alias;
    }
  }
}
}
bgp {
  vpn-apply-export;
  group vpn {
    type internal;
    local-address 10.255.162.91;
    family inet-vpn {
      unicast {
        egress-protection;
      }
    }
    neighbor 10.255.162.84;
    neighbor 10.255.162.89;
  }
}
isis {
  traceoptions {
    file isis.log;
    flag all detail;
  }
  level 2 disable;
  interface all;
  interface fxp0.0 {
    disable;
  }
  interface lo0.0 {
    passive;
  }
}
}

```

```

Device PE3 user@PE3# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.10.40.1/30;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {

```

```
        address 10.10.12.2/30;
    }
    family iso;
    family inet6;
    family mpls;
}
}
lo0 {
    unit 0 {
        family inet {
            address 127.0.0.1/32;
            address 10.255.162.89/32 {
                primary;
            }
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2089.00;
        }
        family inet6 {
            address abcd::10:255:162:89/128 {
                primary;
            }
        }
    }
}
}
```

```
user@PE3# show protocols
rsvp {
    interface all {
        link-protection;
    }
    interface fxp0.0 {
        disable;
    }
}
mpls {
    label-switched-path to_PE2 {
        to 10.255.162.91;
    }
    label-switched-path to_PE1 {
        to 10.255.162.84;
    }
    interface all;
    interface fxp0.0 {
        disable;
    }
    egress-protection {
        context-identifier 1.1.1.1 {
            primary;
            advertise-mode stub-alias;
        }
    }
}
}
bgp {
    vpn-apply-export;
    group vpn {
```

```

    type internal;
    local-address 10.255.162.89;
    family inet-vpn {
        unicast;
    }
    neighbor 10.255.162.84 {
        local-preference 300;
    }
    neighbor 10.255.162.91;
}
}
isis {
    level 2 disable;
    interface all;
    interface fxp0.0 {
        disable;
    }
    interface lo0.0 {
        passive;
    }
}
}

```

Device CE2

```

user@CE2# show interfaces
ge-0/0/0 {
    unit 0 {
        family inet {
            address 10.10.40.2/30;
        }
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 10.10.30.2/30;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 127.0.0.1/32;
            address 10.255.162.88/32 {
                primary;
            }
        }
        family iso {
            address 47.0005.80ff.f800.0000.0108.0001.0102.5516.2088.00;
        }
        family inet6 {
            address abcd::10:255:162:88/128 {
                primary;
            }
        }
    }
}
}

```

Verification

- [Verifying the Routing Instance on page 202](#)
- [Checking the Context Identifier Route on page 208](#)

Verifying the Routing Instance

Purpose Check the routes in the routing table.


```

Action user@PE1> show route 10.10.50 table vpn.inet.0
vpn.inet.0: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.50.0/24      *[BGP/170] 00:01:26, localpref 100, from 10.255.162.96
                   AS path: 65001 I, validation-state: unverified
                   > to 10.10.10.2 via ge-2/0/2.0, Push 16, Push 300064(top)
                   [BGP/170] 00:06:22, localpref 50, from 10.255.162.91
                   AS path: 65001 I, validation-state: unverified
                   > to 10.10.10.2 via ge-2/0/2.0, Push 17, Push 299920(top)

```

```

user@PE1>show route 10.10.50 extensive table vpn.inet.0

```

```

vpn.inet.0: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
10.10.50.0/24 (2 entries, 1 announced)
TSI:
KRT in-kernel 10.10.50.0/24 -> {indirect(1048575)}
Page 0 idx 1, (group vpn type External) Type 1 val 0x9e33490 (adv_entry)
  Advertised metrics:
    Nexthop: Self
    AS path: [65000] 65000 I
    Communities: target:1:1
Path 10.10.50.0 from 10.255.162.96 Vector len 4. Val: 1
  *BGP Preference: 170/-101
    Route Distinguisher: 200:100
    Next hop type: Indirect, Next hop index: 0
    Address: 0x9db63f0
    Next-hop reference count: 6
    Source: 10.255.162.96
    Next hop type: Router, Next hop index: 635
    Next hop: 10.10.10.2 via ge-2/0/2.0, selected
    Label operation: Push 16, Push 300064(top)
    Label TTL action: prop-ttl, prop-ttl(top)
    Load balance label: Label 16: None; Label 300064: None;
    Label element ptr: 0x9db60e0
    Label parent element ptr: 0x9db5e40
    Label element references: 1
    Label element child references: 0
    Label element lsp id: 0
    Session Id: 0x146
    Protocol next hop: 1.1.1.1
    Label operation: Push 16
    Label TTL action: prop-ttl
    Load balance label: Label 16: None;
    Indirect next hop: 0x9e55440 1048575 INH Session ID: 0x14d
    State: < Secondary Active Int Ext ProtectionCand >
    Local AS: 65000 Peer AS: 65000
    Age: 1:28 Metric2: 1
    Validation State: unverified
    Task: BGP_65000.10.255.162.96
    Announcement bits (2): 0-KRT 1-BGP_RT_Background
    AS path: 65001 I
    Communities: target:1:1
    Import Accepted
    VPN Label: 16
    Localpref: 100
    Router ID: 10.255.162.96
    Primary Routing Table bgp.13vpn.0
    Indirect next hops: 1
      Protocol next hop: 1.1.1.1 Metric: 1

```

```

Label operation: Push 16
Label TTL action: prop-ttl
Load balance label: Label 16: None;
Indirect next hop: 0x9e55440 1048575 INH Session ID: 0x14d

Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 10.10.10.2 via ge-2/0/2.0
    Session Id: 0x146
1.1.1.1/32 Originating RIB: inet.3
    Metric: 1    Node path count: 1
    Forwarding nexthops: 1
    Nexthop: 10.10.10.2 via ge-2/0/2.0
BGP    Preference: 170/-51
    Route Distinguisher: 100:100
    Next hop type: Indirect, Next hop index: 0
    Address: 0x9db6390
    Next-hop reference count: 5
    Source: 10.255.162.91
    Next hop type: Router, Next hop index: 636
    Next hop: 10.10.10.2 via ge-2/0/2.0, selected
    Label operation: Push 17, Push 299920(top)
    Label TTL action: prop-ttl, prop-ttl(top)
    Load balance label: Label 17: None; Label 299920: None;
    Label element ptr: 0x9db62c0
    Label parent element ptr: 0x9dc0d00
    Label element references: 1
    Label element child references: 0
    Label element lsp id: 0
    Session Id: 0x146
    Protocol next hop: 10.255.162.91
    Label operation: Push 17
    Label TTL action: prop-ttl
    Load balance label: Label 17: None;
    Indirect next hop: 0x9e55580 1048574 INH Session ID: 0x14c
    State: < Secondary Int Ext ProtectionCand >
    Inactive reason: Local Preference
    Local AS: 65000 Peer AS: 65000
    Age: 6:24 Metric2: 1
    Validation State: unverified
    Task: BGP_65000.10.255.162.91
    AS path: 65001 I
    Communities: target:1:1
    Import Accepted
    VPN Label: 17
    Localpref: 50
    Router ID: 10.255.162.91
    Primary Routing Table bgp.l3vpn.0
    Indirect next hops: 1
        Protocol next hop: 10.255.162.91 Metric: 1
        Label operation: Push 17
        Label TTL action: prop-ttl
        Load balance label: Label 17: None;
        Indirect next hop: 0x9e55580 1048574 INH Session ID: 0x14c

Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 10.10.10.2 via ge-2/0/2.0
    Session Id: 0x146
10.255.162.91/32 Originating RIB: inet.3
    Metric: 1    Node path count: 1

```

```
Forwarding nexthops: 1
Nexthop: 10.10.10.2 via ge-2/0/2.0
```

```
user@PE2> show route table mpls.0
```

```
mpls.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
0          *[MPLS/0] 00:23:33, metric 1
           to table inet.0
0(S=0)     *[MPLS/0] 00:23:33, metric 1
           to table mpls.0
1          *[MPLS/0] 00:23:33, metric 1
           Receive
2          *[MPLS/0] 00:23:33, metric 1
           to table inet6.0
2(S=0)     *[MPLS/0] 00:23:33, metric 1
           to table mpls.0
13         *[MPLS/0] 00:23:33, metric 1
           Receive
17         *[VPN/0] 00:23:33
           to table vpn.inet.0, Pop
299856(S=0) *[MPLS/0] 00:23:33
           to table __1.1.1.1__.mpls.0
299904     *[LDP/9] 00:01:50, metric 1
           > to 10.10.11.2 via xe-8/2/5.0, Pop
299904(S=0) *[LDP/9] 00:01:50, metric 1
           > to 10.10.11.2 via xe-8/2/5.0, Pop
299920     *[LDP/9] 00:01:50, metric 1
           > to 10.10.11.2 via xe-8/2/5.0, Swap 299904
300016     *[LDP/9] 00:01:50, metric 1
           > to 10.10.12.1 via ge-3/0/2.0, Pop
           to table __1.1.1.1__.mpls.0
300016(S=0) *[LDP/9] 00:01:50, metric 1
           > to 10.10.12.1 via ge-3/0/2.0, Pop
           to table __1.1.1.1__.mpls.0
300048     *[LDP/9] 00:01:50, metric 1
           > to 10.10.12.1 via ge-3/0/2.0, Pop
300048(S=0) *[LDP/9] 00:01:50, metric 1
           > to 10.10.12.1 via ge-3/0/2.0, Pop
```

```
user@PE2> show route table __1.1.1.1__.mpls.0
```

```
__1.1.1.1__.mpls.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
16          *[Egress-Protection/170] 00:22:57
           to table __1.1.1.1-vpn__.inet.0
```

```
user@PE2> show route table __1.1.1.1__.mpls.0 extensive
```

```
__1.1.1.1__.mpls.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
16 (1 entry, 1 announced)
```

```
State: < CalcForwarding >
```

```
TSI:
```

```
KRT in-kernel 16 /52 -> {Table}
```

```
*Egress-Protection Preference: 170
```

```
Next table: __1.1.1.1-vpn__.inet.0
```

```
Next-hop index: 649
```

```
Address: 0x9dc2690
```

```
Next-hop reference count: 2
```

```
State: < Active NoReadvrt ForwardingOnly Int Ext >
```

```

Local AS: 65000
Age: 22:59
Validation State: unverified
Task: Protection
Announcement bits (1): 0-KRT
AS path: I
Protecting 2 routes

user@PE2> show route table __1.1.1.1-vpn__.inet.0
__1.1.1.1-vpn__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.30.0/24      *[Egress-Protection/170] 00:02:11
                   to table vpn.inet.0
10.10.50.0/24      *[Egress-Protection/170] 00:02:11
                   > to 10.10.30.2 via ge-3/2/4.0

user@PE2> show route table __1.1.1.1-vpn__.inet.0 extensive
__1.1.1.1-vpn__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.10.30.0/24 (1 entry, 1 announced)
    State: < CalcForwarding >
    TSI:
    KRT in-kernel 10.10.30.0/24 -> {Table}
        *Egress-Protection Preference: 170
        Next table: vpn.inet.0
        Next-hop index: 592
        Address: 0x9dc2630
        Next-hop reference count: 2
        State: < Active NoReadvrt ForwardingOnly Int Ext >
        Local AS: 65000
        Age: 2:13
        Validation State: unverified
        Task: Protection
        Announcement bits (1): 0-KRT
        AS path: I
        Backup route 10.10.30.0 table vpn.inet.0

10.10.50.0/24 (1 entry, 1 announced)
    State: < CalcForwarding >
    TSI:
    KRT in-kernel 10.10.50.0/24 -> {10.10.30.2}
        *Egress-Protection Preference: 170
        Next hop type: Router, Next hop index: 630
        Address: 0x9dc1d90
        Next-hop reference count: 7
        Next hop: 10.10.30.2 via ge-3/2/4.0, selected
        Session Id: 0x147
        State: < Active NoReadvrt ForwardingOnly Int Ext >
        Local AS: 65000
        Age: 2:13
        Validation State: unverified
        Task: Protection
        Announcement bits (1): 0-KRT
        AS path: I
        Backup route 10.10.50.0 table vpn.inet.0

user@PE2> show route table mpls.0 label 17
mpls.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

17                  *[VPN/0] 00:25:06

```

to table vpn.inet.0, Pop

user@PE2> show route table mpls.0 label 17 extensive

mpls.0: 15 destinations, 15 routes (15 active, 0 holddown, 0 hidden)

17 (1 entry, 0 announced)

```
*VPN      Preference: 0
           Next table: vpn.inet.0
           Next-hop index: 0
           Label operation: Pop
           Load balance label: None;
           Label element ptr: 0x9db3920
           Label parent element ptr: 0x0
           Label element references: 1
           Label element child references: 0
           Label element lsp id: 0
           Address: 0x9db3990
           Next-hop reference count: 1
           State: < Active NotInstall Int Ext >
Age: 25:30
           Validation State: unverified
           Task: RT
           AS path: I
```

user@PE3> show route table mpls.0

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

```
0          *[MPLS/0] 00:24:16, metric 1
           to table inet.0
0(S=0)     *[MPLS/0] 00:24:16, metric 1
           to table mpls.0
1          *[MPLS/0] 00:24:16, metric 1
           Receive
2          *[MPLS/0] 00:24:16, metric 1
           to table inet6.0
2(S=0)     *[MPLS/0] 00:24:16, metric 1
           to table mpls.0
13         *[MPLS/0] 00:24:16, metric 1
           Receive
16         *[VPN/0] 00:24:15
           to table vpn.inet.0, Pop
300096     *[LDP/9] 00:02:33, metric 1
           > to 10.10.12.2 via ge-1/1/4.0, Swap 299920
300112     *[LDP/9] 00:02:33, metric 1
           > to 10.10.12.2 via ge-1/1/4.0, Swap 299904
300128     *[LDP/9] 00:02:33, metric 1
           > to 10.10.12.2 via ge-1/1/4.0, Pop
300128(S=0) *[LDP/9] 00:02:33, metric 1
           > to 10.10.12.2 via ge-1/1/4.0, Pop
```

user@PE3> show route table mpls.0 label 16

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

```
16         *[VPN/0] 00:24:22
           to table vpn.inet.0, Pop
```

user@PE3> show route table mpls.0 label 16 extensive

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

16 (1 entry, 0 announced)

```
*VPN      Preference: 0
          Next table: vpn.inet.0
          Next-hop index: 0
          Label operation: Pop
          Load balance label: None;
          Label element ptr: 0x31d1ec0
          Label parent element ptr: 0x0
          Label element references: 1
          Label element child references: 0
          Label element lsp id: 0
          Address: 0x31d1f30
          Next-hop reference count: 1
          State: < Active NotInstall Int Ext >
          Age: 24:24
          Validation State: unverified
          Task: RT
          AS path: I
```

Checking the Context Identifier Route

Purpose Examine the information about the context identifier (1.1.1.1).

```

Action user@PE1> show route 1.1.1.1
inet.0: 47 destinations, 47 routes (46 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[IS-IS/15] 00:04:08, metric 31
                    > to 10.10.10.2 via ge-2/0/2.0

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[LDP/9] 00:04:08, metric 1
                    > to 10.10.10.2 via ge-2/0/2.0, Push 300064

inet.5: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[IS-IS/15] 00:04:08, metric 31, metric2 1
                    > to 10.10.10.2 via ge-2/0/2.0, Push 299856, Push 299920(top)

user@PE2> show route 1.1.1.1
inet.0: 48 destinations, 49 routes (47 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[MPLS/2] 00:26:00, metric 16777215
                    Receive
                    [IS-IS/15] 00:04:17, metric 11
                    > to 10.10.12.1 via ge-3/0/2.0

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[LDP/9] 00:04:17, metric 1
                    > to 10.10.12.1 via ge-3/0/2.0

user@PE2> show mpls context-identifier
ID          Type      Metric    ContextTable
1.1.1.1     protector  16777215  __1.1.1.1__.mpls.0
Total 1, Primary 0, Protector 1

user@PE2> show mpls context-identifier detail
ID: 1.1.1.1
Type: protector, Metric: 16777215, Mode: alias
Context table: __1.1.1.1__.mpls.0, Label out: 299856

Total 1, Primary 0, Protector 1

user@PE3> show route 1.1.1.1
inet.0: 47 destinations, 47 routes (46 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[MPLS/1] 00:26:09, metric 1
                    Receive

inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          *[MPLS/1] 00:26:09, metric 1
                    Receive

```

```
inet.5: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
1.1.1.1/32          *[IS-IS/15] 00:04:27, metric 1, metric2 1
                    > to 10.10.12.2 via ge-1/1/4.0, Push 299856
```

```
user@PE3> show mpls context-identifier
```

ID	Type	Metric	ContextTable
1.1.1.1	primary	1	
Total 1, Primary 1, Protector 0			

```
user@PE3> show mpls context-identifier detail
```

```
ID: 1.1.1.1
  Type: primary, Metric: 1, Mode: alias
```

```
Total 1, Primary 1, Protector 0
```

Related Documentation

- *Egress Protection for Layer 3 VPN Edge Protection Overview*

CHAPTER 4

Configuring MPLS-Signaled LSPs

- [Configuring the Ingress and Egress Router Addresses for LSPs on page 212](#)
- [Configuring Primary and Secondary LSPs on page 214](#)
- [Configuring a Text Description for LSPs on page 217](#)
- [Configuring the Entropy Label for LSPs on page 218](#)
- [Configuring Corouted Bidirectional LSPs on page 220](#)
- [Configuring Ultimate-Hop Popping for LSPs on page 222](#)
- [Configuring an LSP Across ASs on page 225](#)
- [Configuring Fast Reroute on page 226](#)
- [Configuring the Optimization Interval for Fast Reroute Paths on page 228](#)
- [Adding LSP-Related Routes to the inet.3 or inet6.3 Routing Table on page 228](#)
- [Configuring the Connection Between Ingress and Egress Routers on page 229](#)
- [Configuring LSP Metrics on page 230](#)
- [Configuring CSPF Tie Breaking on page 232](#)
- [Configuring Load Balancing Based on MPLS Labels on DPC I-Chip-Based Hardware on page 232](#)
- [Disabling Normal TTL Decrementing on page 236](#)
- [Configuring MPLS Soft Preemption on page 238](#)
- [Disabling Constrained-Path LSP Computation on page 239](#)
- [Configuring Administrative Groups for LSPs on page 240](#)
- [Configuring Extended Administrative Groups for LSPs on page 242](#)
- [Configuring Preference Values for LSPs on page 243](#)
- [Disabling Path Route Recording by LSPs on page 244](#)
- [Configuring Class of Service for MPLS LSPs on page 244](#)
- [Achieving a Make-Before-Break, Hitless Switchover for LSPs on page 246](#)
- [Configuring Adaptive LSPs on page 249](#)
- [Configuring Priority and Preemption for LSPs on page 250](#)
- [Optimizing Signaled LSPs on page 251](#)
- [Configuring the Smart Optimize Timer for LSPs on page 255](#)

- [Limiting the Number of Hops in LSPs on page 256](#)
- [Configuring the Bandwidth Value for LSPs on page 256](#)
- [Automatic Bandwidth Allocation for LSPs on page 257](#)
- [Configuring Automatic Bandwidth Allocation for LSPs on page 257](#)
- [Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs on page 264](#)
- [Configuring Hot Standby of Secondary Paths for LSPs on page 267](#)
- [Damping Advertisement of LSP State Changes on page 268](#)

Configuring the Ingress and Egress Router Addresses for LSPs

The following sections describe how to specify the addresses of an LSP's ingress and egress routers:

- [Configuring the Ingress Router Address for LSPs on page 212](#)
- [Configuring the Egress Router Address for LSPs on page 212](#)
- [Preventing the Addition of Egress Router Addresses to Routing Tables on page 213](#)

Configuring the Ingress Router Address for LSPs

The local router always is considered to be the ingress router, which is the beginning of the LSP. The software automatically determines the proper outgoing interface and IP address to use to reach the next router in an LSP.

By default, the router ID is chosen as the address of the ingress router. To override the automatic selection of the source address, specify a source address in the **from** statement:

from *address*;

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls label-switched-path *lsp-name*]**
- **[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]**

The outgoing interface used by the LSP is not affected by the source address that you configure.

Configuring the Egress Router Address for LSPs

When configuring an LSP, you must specify the address of the egress router by including the **to** statement:

to *address*;

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls label-switched-path *lsp-name*]**
- **[edit protocols mpls static-label-switched-path *lsp-name*]**

- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **static-label-switched-path** *lsp-name*]

When you are setting up a signaled LSP, the **to** statement is the only required statement. All other statements are optional.

After the LSP is established, the address of the egress router is installed as a host route in the routing table. This route can then be used by BGP to forward traffic.

To have the software send BGP traffic over an LSP, the address of the egress router is the same as the address of the BGP next hop. You can specify the egress router's address as any one of the router's interface addresses or as the BGP router ID. If you specify a different address, even if the address is on the same router, BGP traffic is not sent over the LSP.

To determine the address of the BGP next hop, use the **show route detail** command. To determine the destination address of an LSP, use the **show mpls lsp** command. To determine whether a route has gone through an LSP, use the **show route** or **show route forwarding-table** command. In the output of these last two commands, the **label-switched-path** or **push** keyword included with the route indicates it has passed through an LSP. Also, use the **traceroute** command to trace the actual path to which the route leads. This is another indication whether a route has passed through an LSP.

You also can manipulate the address of the BGP next hop by defining a BGP import policy filter that sets the route's next-hop address.

Preventing the Addition of Egress Router Addresses to Routing Tables

You must configure an address using the **to** statement for all LSPs. This address is always installed as a /32 prefix in the inet.3 or inet.0 routing tables. You can prevent the egress router address configured using the **to** statement from being added to the inet.3 and inet.0 routing tables by including the **no-install-to-address** statement.

Some reasons not to install the **to** statement address in the inet.3 and inet.0 routing tables include the following:

- Allow Constrained Shortest Path First (CSPF) RSVP LSPs to be mapped to traffic intended for secondary loopback addresses. If you configure an RSVP tunnel, including the **no-install-to-address** statement, and then configure an **install pfx/ <active>** policy later, you can do the following:
 - Verify that the LSP was set up correctly without impacting traffic.
 - Map traffic to the LSP in incremental steps.
 - Map traffic to the destination loopback address (the BGP next hop) by removing the **no-install-to-address** statement once troubleshooting is complete.
- Prevent CCC connections from losing IP traffic. When an LSP determines that it does not belong to a connection, it installs the address specified with the **to** statement in the inet.3 routing table. IP traffic is then forwarded to the CCC remote endpoint, which can cause some types of PICs to fail.

To prevent the egress router address configured using the **to** statement from being added to the inet.3 and inet.0 routing tables, include the **no-install-to-address** statement:

```
no-install-to-address;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls [label-switched-path](#) *lsp-name*]
- [edit protocols mpls [static-label-switched-path](#) *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls [label-switched-path](#) *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls [static-label-switched-path](#) *lsp-name*]

Configuring Primary and Secondary LSPs

By default, an LSP routes itself hop-by-hop toward the egress router. The LSP tends to follow the shortest path as dictated by the local routing table, usually taking the same path as destination-based, best-effort traffic. These paths are “soft” in nature because they automatically reroute themselves whenever a change occurs in a routing table or in the status of a node or link.

To configure the path so that it follows a particular route, create a named path using the **path** statement, as described in [“Creating Named Paths” on page 60](#). Then apply the named path by including the **primary** or **secondary** statement. A named path can be referenced by any number of LSPs.

To configure primary and secondary paths for an LSP, complete the steps in the following sections:

- [Configuring Primary and Secondary Paths for an LSP on page 214](#)
- [Configuring the Revert Timer for LSPs on page 215](#)
- [Specifying the Conditions for Path Selection on page 216](#)

Configuring Primary and Secondary Paths for an LSP

The **primary** statement creates the primary path, which is the LSP’s preferred path. The **secondary** statement creates an alternative path. If the primary path can no longer reach the egress router, the alternative path is used.

To configure primary and secondary paths, include the **primary** and **secondary** statements:

```
primary path-name {  
  ...  
}  
secondary path-name {  
  ...  
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols mpls [label-switched-path](#) *lsp-name*]

- `[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name]`

When the software switches from the primary to a secondary path, it continuously attempts to revert to the primary path, switching back to it when it is again reachable, but no sooner than the retry time specified in the **retry-timer** statement. (For more information, see “Configuring the Connection Between Ingress and Egress Routers” on page 229.)

You can configure zero or one primary path. If you do not configure a primary path, the first secondary path that is established is selected as the path.

You can configure zero or more secondary paths. All secondary paths are equal, and the software tries them in the order that they are listed in the configuration. The software does not attempt to switch among secondary paths. If the current secondary path is not available, the next one is tried. To create a set of equal paths, specify secondary paths without specifying a primary path.

If you do not specify any named paths, or if the path that you specify is empty, the software makes all routing decisions necessary to reach the egress router.

Configuring the Revert Timer for LSPs

For LSPs configured with both primary and secondary paths, it is possible to configure the revert timer. If a primary path goes down and traffic is switched to the secondary path, the revert timer specifies the amount of time (in seconds) that the LSP must wait before it can revert traffic back to a primary path. If during this time, the primary path experiences any connectivity problems or stability problems, the timer is restarted. You can configure the revert timer for both static and dynamic LSPs.

The Junos OS also makes a determination as to which path is the preferred path. The preferred path is the path that has not encountered any difficulty in the last revert timer period. If both the primary and secondary paths have encountered difficulty, neither path is considered preferred. However, if one of the paths is dynamic and the other static, the dynamic path is selected as the preferred path.

If you have configured BFD on the LSP, Junos OS waits until the BFD session comes up on the primary path before starting the revert timer counter.

The range of values you can configure for the revert timer is 0 through 65,535 seconds. The default value is 60 seconds.

If you configure a value of 0 seconds, the traffic on the LSP, once switched from the primary path to the secondary path, remains on the secondary path permanently (until the network operator intervenes or until the secondary path goes down).

You can configure the revert timer for all LSPs on the router at the `[edit protocols mpls]` hierarchy level or for a specific LSP at the `[edit protocols mpls label-switched-path lsp-name]` hierarchy level.

To configure the revert timer, include the **revert-timer** statement:

```
revert-timer seconds;
```

For a list of hierarchy levels at which you can include this statement, see the summary section for this statement.

Specifying the Conditions for Path Selection

When you have configured both primary and secondary paths for an LSP, you may need to ensure that only a specific path is used.

The **select** statement is optional. If you do not include it, MPLS uses an automatic path selection algorithm.

The **manual** and **unconditional** options do the following:

- **manual**—The path is immediately selected for carrying traffic as long as it is up and stable. Traffic is sent to other working paths if the current path is down or degraded (receiving errors). This parameter overrides all other path attributes except the **select unconditional** statement.
- **unconditional**—The path is selected for carrying traffic unconditionally, regardless of whether the path is currently down or degraded (receiving errors). This parameter overrides all other path attributes.

Because the **unconditional** option switches to a path without regard to its current status, be aware of the following potential consequences of specifying it:

- If a path is not currently up when you enable the **unconditional** option, traffic can be disrupted. Ensure that the path is functional before specifying the **unconditional** option.
- Once a path is selected because it has the **unconditional** option enabled, all other paths for the LSP are gradually cleared, including the primary and standby paths. No path can act as a standby to an unconditional path, so signaling those paths serves no purpose.

For a specific path, the **manual** and **unconditional** options are mutually exclusive. You can include the **select** statement with the **manual** option in the configuration of only one of an LSP's paths, and the **select** statement with the **unconditional** option in the configuration of only one other of its paths.

Enabling or disabling the **manual** and **unconditional** options for the **select** statement while LSPs and their paths are up does not disrupt traffic.

To specify that a path be selected for carrying traffic if it is up and stable for at least the revert timer window, include the **select** statement with the **manual** option:

```
select manual;
```

To specify that a path should always be selected for carrying traffic, even if it is currently down or degraded, include the **select** statement with the **unconditional** option:

```
select unconditional;
```

You can include the **select** statement at the following hierarchy levels:

- `[edit protocols mpls label-switched-path lsp-name (primary | secondary) path-name]`

- [edit logical-systems *logical-system-name* protocols mpls *label-switched-path* *lsp-name* (*primary* | *secondary*) *path-name*]

Configuring a Text Description for LSPs

You can provide a textual description for an LSP by enclosing any descriptive text that includes spaces within quotation marks (" "). The descriptive text you include is displayed in the detail output of the **show mpls lsp** or the **show mpls container-lsp** command.

Adding a text description for an LSP has no effect on the operation of the LSP. The LSP text description can be no more than 80 characters in length.

To provide a textual description for an LSP, include the **description** statement at any of the following hierarchy levels:

- [edit protocols mpls *label-switched-path* *lsp-name*]
- [edit protocols mpls *container-label-switched-path* *lsp-name*]
- [edit protocols mpls *static-label-switched-path* *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls *label-switched-path* *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls *static-label-switched-path* *lsp-name*]

Before you begin:

- Configure the device interfaces.
- Configure the device for network communication.
- Enable MPLS on the device interfaces.
- Configure an LSP in the MPLS domain.

To add a text description for an LSP:

1. Enter any text describing the LSP.

```
[edit protocols mpls lsp lsp-name]  
user@host# set description text
```

For example:

```
[edit protocols mpls lsp LSP1]  
user@host# set description "Connecting remote device"
```

2. Verify and commit the configuration.

For example:

```
[edit protocols mpls lsp]  
user@host# set protocols mpls label-switched-path LSP1 to 1.1.1.1  
user@host# set protocols mpls label-switched-path LSP1 description "Connecting  
remote device"  
user@host# set protocols mpls interface ge-1/0/8.0  
[edit]
```

```
user@host# commit
commit complete
```

3. View the description of an LSP using the **show mpls lsp detail** or **show mpls container-lsp detail** command, depending on the type of LSP configured.

```
user@host> show mpls lsp detail
```

```
Ingress LSP: 1 sessions
```

```
1.1.1.1
```

```
From: 0.0.0.0, State: Up, ActiveRoute: 1, LSPname: LSP1
```

```
Description: Connecting remote device
```

```
ActivePath: (none)
```

```
LSPtype: Static Configured, Penultimate hop popping
```

```
LoadBalance: Random
```

```
Encoding type: Packet, Switching type: Packet, GPID: IPv4
```

```
Primary State: Up
```

```
Priorities: 7 0
```

```
SmartOptimizeTimer: 180
```

```
No computed ERO.
```

```
Total 1 displayed, Up 1, Down 0
```

```
Egress LSP: 0 sessions
```

```
Total 0 displayed, Up 0, Down 0
```

```
Transit LSP: 0 sessions
```

```
Total 0 displayed, Up 0, Down 0
```

Related Documentation

- [label-switched-path on page 884](#)
- [Configuring LSP Metrics on page 230](#)
- [Minimum MPLS Configuration on page 59](#)

Configuring the Entropy Label for LSPs

The insertion of entropy labels for an LSP enables transit routers to load-balance MPLS traffic across ECMP paths or Link Aggregation groups using just the MPLS label stack as a hash input without having to rely on deep packet inspection. Deep packet inspection requires more of the router's processing power and different routers have differing deep-packet inspection capabilities.

To configure the entropy label for an LSP, complete the following steps:

1. On the ingress router, include the **entropy-label** statement at the **[edit protocols mpls labeled-switched-path *labeled-switched-path-name*]** hierarchy level or at the **[edit protocols mpls static-labeled-switched-path *labeled-switched-path-name* ingress]** hierarchy level. The entropy label is added to the MPLS label stack and can be processed in the forwarding plane.

```
entropy-label;
```



NOTE: This is only applicable for RSVP and static LSPs.

2. On the ingress router, you can configure an ingress policy for LDP-signaled LSPs:

```
entropy-label {
  ingress-policy policy-name;
}
```

Configure the ingress policy at the **[edit policy-options]** hierarchy level:

```
policy-options {
  policy-statement policy-name {
    term term-name {
      from {
        prefix-list prefix-list-name;
      }
      then actions;
    }
  }
}
```

The following shows an example of an entropy label ingress policy.

```
policy-options {
  policy-statement entropy-policy {
    term no-insert-entropy-label {
      from {
        prefix-list no-entropy-label-fec;
      }
      then accept;
    }
  }
}
```

3. (Optional) By default, routers that support the pushing and popping of entropy labels are configured with the **load-balance-label-capability** statement at the **[edit forwarding-options]** hierarchy level to signal the labels on a per-LSP basis. If the peer router is not equipped to handle load-balancing labels, you can prevent the provider edge (PE) router from signaling the entropy label capability by configuring the **no-load-balance-label-capability** statement at the **[edit forwarding-options]** hierarchy level.

```
[edit forwarding-options]
user@PE no-load-balance-label-capability;
```

Transit routers require no configuration. The presence of the entropy label indicates to the transit router to load balance based solely on the MPLS label stack.

Penultimate hop routers pop the entropy label by default.

Related Documentation


- *Configuring Per-Packet Load Balancing*
- *enhanced-hash-key*
- [entropy-label on page 860](#)
- *hash-key*
- [ingress-policy on page 881](#)

Configuring Corouted Bidirectional LSPs

A corouted bidirectional packet LSP is a combination of two LSPs sharing the same path between a pair of ingress and egress nodes, as shown in [Figure 27 on page 220](#). It is established using the GMPLS extensions to RSVP-TE. This type of LSP can be used to carry any of the standard types of MPLS-based traffic, including Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs. You can configure a single BFD session for the bidirectional LSP (you do not need to configure a BFD session for each LSP in each direction). You can also configure a single standby bidirectional LSP to provide a backup for the primary bidirectional LSP. Corouted bidirectional LSPs are supported for both penultimate hop popping (PHP) and ultimate hop popping (UHP).

High availability is available for bidirectional LSPs. You can enable graceful restart and nonstop active routing. Graceful restart and nonstop active routing are supported when the restarting router is the ingress, egress, or transit router for the bidirectional LSP.

Figure 27: Corouted Bidirectional LSP



To configure a corouted bidirectional LSP:

1. In configuration mode, configure the ingress router for the LSP and include the **corouted-bidirectional** statement to specify that the LSP be established as a corouted bidirectional LSP.

The path is computed using CSPF and initiated using RSVP signaling (just like a unidirectional RSVP signaled LSP). Both the path to the egress router and the reverse path from the egress router are created when this configuration is committed.

```
[edit protocols mpls]
user@PE1# set label-switched-path sample-lsp corouted-bidirectional
```

2. (Optional) For a reverse path, configure an LSP on the egress router and include the **corouted-bidirectional-passive** statement to associate the LSP with another LSP.

No path computation or signaling is used for this LSP since it relies on the path computation and signaling provided by the ingress LSP. You cannot configure both the **corouted-bidirectional** statement and the **corouted-bidirectional-passive** statement on the same LSP.

```
[edit protocols mpls]
user@PE1# set label-switched-path sample-lsp-reverse-path
corouted-bidirectional-passive
```

This statement also makes it easier to debug corouted bidirectional LSPs. If you configure the **corouted-bidirectional-passive** statement (again, on the egress router), you can issue **ping mpls lsp-end-point**, **ping mpls ldp**, **ping mpls rsvp**, **traceroute mpls**

`ldp`, and `traceroute mpls rsvp` commands to test the core-routed bidirectional LSP from the egress router.

3. Use the `show mpls lsp extensive` and the `show rsvp session extensive` commands to display information about the bidirectional LSP.

The following shows output for the `show rsvp session extensive` command when run on an ingress router with a bidirectional LSP configured:

```
user@PE1> show rsvp session extensive
Ingress RSVP: 2 sessions

10.255.14.39
  From: 10.255.14.43, LSPstate: Up, ActiveRoute: 0
  LSPname: l-to-h, LSPpath: Primary
  LSPtype: Static Configured
  Bidirectional, Upstream label in: 3, Upstream label out: -
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 300032
  Resv style: 1 FF, Label in: -, Label out: 300032
  Time left: -, Since: Tue May 31 08:49:25 2011
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 24617 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.1.1.2 (ge-0/0/0.0) 3396 pkts
  RESV rcvfrom: 10.1.1.2 (ge-0/0/0.0) 3394 pkts
  PATH notifyto: localclient
  RESV notifyto: 10.255.14.39
  Protection attributes: primary, working, 1:N protection
  Association attributes: recovery, src 10.255.14.43, id 1
  Explct route: 10.1.1.2 10.1.2.2 10.1.3.2
  Record route: 10.1.1.2 10.1.2.2 10.1.3.2

10.255.14.39
  From: 10.255.14.43, LSPstate: Up, ActiveRoute: 0
  LSPname: l-to-h, LSPpath: Secondary
  LSPtype: Static Configured
  Bidirectional, Upstream label in: 3, Upstream label out: -
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 300032
  Resv style: 1 FF, Label in: -, Label out: 300032
  Time left: -, Since: Tue May 31 08:49:25 2011
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 24617 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.1.1.2 (ge-0/0/0.0) 3396 pkts
  RESV rcvfrom: 10.1.1.2 (ge-0/0/0.0) 3394 pkts
  Protection attributes: primary, protecting
  Association attributes: recovery, src 10.255.14.43, id 1
  Explct route: 10.2.1.2 10.2.2.2 10.2.3.2
  Record route: 10.2.1.2 10.2.2.2 10.2.3.2
Total 2 displayed, Up 2, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0


- Related Documentation**
- [Configuring Ultimate-Hop Popping for LSPs on page 222](#)
 - [Configuring LDP Graceful Restart on page 533](#)
 - [Configuring RSVP Graceful Restart on page 514](#)
 - [Configuring Nonstop Active Routing](#)

Configuring Ultimate-Hop Popping for LSPs

By default, RSVP-signaled LSPs use penultimate-hop popping (PHP).


[Figure 28 on page 222](#) illustrates a penultimate-hop popping LSP between Router PE1 and Router PE2. Router CE1 forwards a packet to its next hop (Router PE1), which is also the LSP ingress. Router PE1 pushes label 1 on the packet and forwards the labeled packet to Router P1. Router P1 completes the standard MPLS label swapping operation, swapping label 1 for label 2, and forwards the packet to Router P2. Since Router P2 is the penultimate-hop router for the LSP to Router PE2, it first pops the label and then forwards the packet to Router PE2. When Router PE2 receives it, the packet can have a service label, an explicit-null label, or just be a plain IP or VPLS packet. Router PE2 forwards the unlabeled packet to Router CE2.

Figure 28: Penultimate-Hop Popping for an LSP



You can also configure ultimate-hop popping (UHP) (as shown in [Figure 29 on page 222](#)) for RSVP-signaled LSPs. Some network applications can require that packets arrive at the egress router (Router PE2) with a non-null outer label. For an ultimate-hop popping LSP, the penultimate router (Router P2 in [Figure 29 on page 222](#)) performs the standard MPLS label swapping operation (in this example, label 2 for label 3) before forwarding the packet to egress Router PE2. Router PE2 pops the outer label and performs a second lookup of the packet address to determine the end destination. It then forwards the packet to the appropriate destination (either Router CE2 or Router CE4).

Figure 29: Ultimate-Hop Popping for an LSP



The following network applications require that you configure UHP LSPs:

- MPLS-TP for performance monitoring and in-band OAM
- Edge protection virtual circuits

The following features do not support the UHP behavior:

- LDP-signaled LSPs
- Static LSPs
- Point-to-multipoint LSPs
- CCC
- **traceroute** command

For more information about UHP behavior, see Internet draft [draft-ietf-mpls-rsvp-te-no-php-oob-mapping-01.txt](#), *Non PHP behavior and Out-of-Band Mapping for RSVP-TE LSPs*.

For point-to-point RSVP-signaled LSPs, UHP behavior is signaled from the LSP ingress. Based on the ingress router configuration, RSVP can signal the UHP LSP with the non-PHP flag set. RSVP PATH messages carry the two flags in the LSP-ATTRIBUTES object. When the egress router receives the PATH message, it assigns a non-null label to the LSP. RSVP also creates and installs two routes in the mpls.0 routing table. S refers to the S bit of the MPLS label, which indicates whether or not the bottom of the label stack has been reached.

- Route S=0—Indicates that there are more labels in the stack. The next hop for this route points to the mpls.0 routing table, triggering a chained MPLS label lookup to discover the remaining MPLS labels in the stack.
- Route S=1—Indicates that there are no more labels. The next hop points to the inet.0 routing table if the platform supports chained and multi-family lookup. Alternatively, the label route can point to a VT interface to initiate IP forwarding.

If you enable UHP LSPs, MPLS applications such as Layer 3 VPNs, VPLS, Layer 2 VPNs, and Layer 2 circuits can use the UHP LSPs. The following explains how UHP LSPs affect the different types of MPLS applications:

- Layer 2 VPNs and Layer 2 circuits—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label (S=0) is the UHP label, and the inner label (S=1) is the VC label. A lookup based on the transport label results in a table handle for the mpls.0 routing table. There is an additional route in the mpls.0 routing table corresponding to the inner label. A lookup based on the inner label results in the CE router next hop.
- Layer 3 VPN—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label (S=0) is the UHP label, and the inner label is the VPN label (S=1). A lookup based on the transport label results in the table handle for the mpls.0 routing table. There are two cases in this scenario. By default, Layer 3 VPNs advertise the per-next hop label. A lookup based on the inner label results in the next hop toward

the CE router. However, if you have configured the **vrf-table-label** statement for the Layer 3 VPN routing instance, the inner LSI label points to the VRF routing table. An IP lookup is also completed for the VRF routing table.



NOTE: UHP for Layer 3 VPNs configured with the **vrf-table-label** statement is supported on MX 3D Universal Edge Routers only.

- **VPLS**—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label is the transport label (S=0) and the inner label is the VPLS label (S=1). A lookup based on the transport label results in the table handle for the mpls.0 routing table. A lookup based on the inner label in mpls.0 routing table results in the LSI tunnel interface of the VPLS routing instance if tunnel-services is not configured (or a VT interface not available). MX 3D Series routers support chained lookup and multi-family lookup.



NOTE: UHP for VPLS configured with the **no-tunnel-service** statement is supported on MX 3D Series routers only.

- **IPv4 over MPLS**—A packet arrives at the PE router (egress of the UHP LSP) with one label (S=1). A lookup based on this label returns a VT tunnel interface. Another IP lookup is completed on the VT interface to determine where to forward the packet. If the routing platform supports multi-family and chained lookups (for example, MX 3D routers and PTX Series Packet Transport Routers), lookup based on label route (S=1) points to the inet.0 routing table.
- **IPv6 over MPLS**—For IPv6 tunneling over MPLS, PE routers advertise IPv6 routes to each other with a label value of 2. This is the explicit null label for IPv6. As a result, the forwarding next hops for IPv6 routes that are learned from remote PE routers normally push two labels. The inner label is 2 (it could be different if the advertising PE router is from another vendor), and the router label is the LSP label. Packets arrive at the PE router (egress of the UHP LSP) with two labels. The outer label is the transport label (S=0), and the inner label is the IPv6 explicit-null label (label 2). Lookup based on the inner label in the mpls.0 routing table redirects back to the mpls.0 routing table. On MX 3D Series routers, the inner label (label 2) is stripped off and an IPv6 lookup is done using the inet6.0 routing table.
- **Enabling both PHP and UHP LSPs**—You can configure both PHP and UHP LSPs over the same network paths. You can separate PHP and UHP traffic by selecting forwarding LSP next hops using a regular expression with the **install-nexthop** statement. You can also separate traffic by simply naming the LSPs appropriately.

The following statements enable ultimate-hop popping for an LSP. You can enable this feature on a specific LSP or for all of the ingress LSPs configured on the router. Configure these statements on the router at the LSP ingress.

1. To enable ultimate-hop popping, include the **ultimate-hop-popping** statement:

```
ultimate-hop-popping;
```

Include this statement at the **[edit protocols mpls label-switched-path *label-switched-path-name*]** hierarchy level to enable ultimate-hop popping on a specific LSP. Include this statement at the **[edit protocols mpls]** hierarchy level to enable ultimate-hop popping on all of the ingress LSPs configured on the router. You can also configure the **ultimate-hop-popping** statement under the equivalent **[edit logical-routers]** hierarchy levels.



NOTE: When you enable ultimate-hop popping, RSVP attempts to resignal existing LSPs as ultimate-hop popping LSPs in a make-before-break fashion. If an egress router does not support ultimate-hop popping, the existing LSP is torn down (RSVP sends a PathTear message along an LSP's path, removing the path state and dependent reservation state and releasing the associated networking resources).

If you disable ultimate-hop popping, RSVP resignals existing LSPs as penultimate-hop popping LSPs in a make-before-break fashion.

2. If you want to enable both ultimate-hop-popping and chained next hops on MX 3D Series routers only, you also need to configure the **enhanced-ip** option for the **network-services** statement:

```
network-services enhanced-ip;
```

You configure this statement at the **[edit chassis]** hierarchy level. Once you have configured the **network-services** statement, you need to reboot the router to enable UHP behavior.

Related Documentation

- [MPLS Label Allocation on page 26](#)
- [Configuring Corouted Bidirectional LSPs on page 220](#)
- *network-services*
- [ultimate-hop-popping on page 965](#)

Configuring an LSP Across ASs

You can configure an LSP to traverse multiple areas in a network by including the **inter-domain** statement as a part of the LSP configuration. This statement allows the router to search for routes in the IGP database. You need to configure this statement on routers that might be unable to locate a path using intra-domain CSPF (by looking in the traffic engineering database (TED)). When you configure inter-area LSPs, the **inter-domain** statement is required.

Before you begin:

- Configure the device interfaces with family MPLS.
- Configure the device router ID and autonomous system number.
- Enable MPLS and RSVP on the router and transit interfaces.

- Configure your IGP to support traffic engineering.
- Set up an LSP from the ingress to the egress router.

To configure an LSP across multiple ASs on the ingress label-switched router (LER):

1. Enable MPLS on all the interfaces (excluding the management interface).

```
[edit protocols]
user@LER# set mpls interface all
user@LER# set mpls interface fxp0.0 disable
```

2. Enable RSVP on all the interfaces (excluding the management interface).

```
[edit protocols]
user@LER# set rsvp interface all
user@LER# set rsvp interface fxp0.0 disable
```

3. Configure the inter-area LSP.

```
[edit protocols]
user@LER# set mpls label-switched-path inter-area-LSP-name to
  egress-LER-ip-address
user@LER# set mpls label-switched-path inter-area-LSP-name inter-domain
```

4. Verify and commit the configuration.

```
[edit protocols]
user@LER# set rsvp interface ge-0/0/0.0
user@LER# set rsvp interface lo0.0
user@LER# set rsvp interface fxp0.0 disable
user@LER# set mpls statistics traffic-class-statistics
user@LER# set mpls label-switched-path R1-R2 to 20.0.0.1
user@LER# set mpls label-switched-path R1-R2 inter-domain
user@LER# set mpls interface ge-0/0/0.0
user@LER# set mpls interface lo0.0
user@LER# set mpls interface fxp0.0 disable
user@LER# set ospf traffic-engineering
user@LER# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@LER# set ospf area 0.0.0.0 interface lo0.0
```

Related Documentation

- [inter-domain on page 883](#)

Configuring Fast Reroute

Fast reroute provides a mechanism for automatically rerouting traffic on an LSP if a node or link in an LSP fails, thus reducing the loss of packets traveling over the LSP.

To configure fast reroute on an LSP, include the **fast-reroute** statement on the ingress router (or switch):

```
fast-reroute {
  (bandwidth bps | bandwidth-percent percentage);
  (exclude [ group-names ] | no-exclude );
  hop-limit number;
  (include-all [ group-names ] | no-include-all);
```



```
(include-any [ group-names ] | no-include-any);
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

You do not need to configure fast reroute on the LSP's transit and egress routers (or switches). Once fast reroute is enabled, the ingress router (or switch) signals all the downstream routers (or switches) that fast reroute is enabled on the LSP, and each downstream router does its best to set up detours for the LSP. If a downstream router does not support fast reroute, it ignores the request to set up detours and continues to support the LSP. A router that does not support fast reroute will cause some of the detours to fail, but otherwise has no impact on the LSP.



NOTE: To enable PFE fast reroute, configure a routing policy statement with the **load-balance per-packet** statement at the [edit policy-options policy-statement *policy-name* then] hierarchy level on each of the routers where traffic might be rerouted. See also “[Configuring Load Balancing Across RSVP LSPs](#)” on page 481.

By default, no bandwidth is reserved for the rerouted path. To allocate bandwidth for the rerouted path, include either the **bandwidth** statement or the **bandwidth-percent** statement. You can only include one of these statements at a time. If you do not include either the **bandwidth** statement or the **bandwidth-percent** statement, the default setting is to not reserve bandwidth for the detour path.

When you include the **bandwidth** statement, you can specify the specific amount of bandwidth (in bits per second [bps]) you want to reserve for the detour path. The bandwidth does not need to be identical to that allocated for the LSP.

When you specify a bandwidth percent using the **bandwidth-percent** statement, the detour path bandwidth is computed by multiplying the bandwidth percentage by the bandwidth configured for the main traffic-engineered LSP. For information about how to configure the bandwidth for a traffic-engineered LSP, see “[Configuring Traffic-Engineered LSPs](#)” on page 326.

Hop-limit constraints define how many more routers a detour is allowed to traverse compared with the LSP itself. By default, the hop limit is set to 6. For example, if an LSP traverses 4 routers, any detour for the LSP can be up to 10 (that is, 4 + 6) router hops, including the ingress and egress routers.

By default, a detour inherits the same administrative (coloring) group constraints as its parent LSP when CSPF is determining the alternate path. Administrative groups, also known as link coloring or resource class, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. If you specify the **include-any** statement when configuring the parent LSP, all links traversed by the alternate session must have at least one color found in the list of groups.

If you specify the **include-all** statement when configuring the parent LSP, all links traversed by the alternate session must have all of the colors found in the list of groups. If you specify the **exclude** statement when configuring the parent LSP, none of the links must have a color found in the list of groups. For more information about administrative group constraints, see [“Configuring Administrative Groups for LSPs” on page 240](#).

**Related
Documentation**

- [Fast Reroute Overview on page 46](#)
- *MPLS Feature Support on QFX Series and EX4600 Switches*

Configuring the Optimization Interval for Fast Reroute Paths

You can enable path optimization for fast reroute by configuring the fast reroute optimize timer. The optimize timer triggers a periodic optimization process that recomputes the fast reroute detour LSPs to use network resources more efficiently.

To enable fast reroute path optimization, specify the number of seconds using the `optimize-timer` option for the **fast-reroute** statement:

```
fast-reroute seconds;
```

You can include this statement at the following hierarchy levels:

- `[edit protocols rsvp]`
- `[edit logical-systems logical-system-name protocols rsvp]`

Adding LSP-Related Routes to the inet.3 or inet6.3 Routing Table

By default, a host route toward the egress router is installed in the `inet.3` or `inet6.3` routing table. (The host route address is the one you configure in the **to** statement.) Installing the host route allows BGP to perform next-hop resolution. It also prevents the host route from interfering with prefixes learned from dynamic routing protocols and stored in the `inet.0` or `inet6.0` routing table.

Unlike the routes in the `inet.0` or `inet6.0` table, routes in the `inet.3` or `inet6.3` table are not copied to the Packet Forwarding Engine, and hence they cause no changes in the system forwarding table directly. You cannot use the **ping** or **traceroute** command through these routes. The only use for `inet.3` or `inet6.3` is to permit BGP to perform next-hop resolution. To examine the `inet.3` or `inet6.3` table, use the **show route table inet.3** or **show route table inet6.3** command.

To inject additional routes into the `inet.3` or `inet6.3` routing table, include the **install** statement:

```
install {  
    destination-prefix <active>;  
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit protocols mpls **static-label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **static-label-switched-path** *lsp-name*]

The specified routes are installed as aliases into the routing table when the LSP is established. Installing additional routes allows BGP to resolve next hops within the specified prefix and to direct additional traffic for these next hops to a particular LSP.

Including the **active** option with the **install** statement installs the specified prefix into the inet.0 or inet6.0 routing table, which is the primary forwarding table. The result is a route that is installed in the forwarding table any time the LSP is established, which means you can ping or trace the route. Use this option with care, because this type of prefix is very similar to a static route.

You use alias routes for routers that have multiple addresses being used as BGP next hops, or for routers that are not MPLS capable. In either of these cases, the LSP can be configured to another MPLS capable system within the local domain, which then acts as a “border” router. The LSP then terminates on the border router and, from that router, Layer 3 forwarding takes the packet to the true next-hop router.

In the case of an interconnect, the domain's border router can act as the proxy router and can advertise the prefix for the interconnect if the border router is not setting the BGP next hop to itself.

In the case of a point of presence (POP) that has routers that do not support MPLS, one router (for example, a core router) that supports MPLS can act as a proxy for the entire POP and can inject a set of prefixes that cover the POP. Thus, all routers within the POP can advertise themselves as interior BGP (IBGP) next hops, and traffic can follow the LSP to reach the core router. This means that normal IGP routing would prevail within the POP.

You cannot use the **ping** or **traceroute** commands on routes in the inet.3 or inet6.3 routing table.

For BGP next-hop resolution, it makes no difference whether a route is in inet.0/inet6.0 or inet.3/inet6.3; the route with the best match (longest mask) is chosen. Among multiple best-match routes, the one with the highest preference value is chosen.

Configuring the Connection Between Ingress and Egress Routers

The ingress router might make many attempts to connect and reconnect to the egress router using the primary path. You can control how often the ingress router tries to establish a connection using the primary path and how long it waits between retry attempts.

The retry timer configures how long the ingress router waits before trying to connect again to the egress router using the primary path. The default retry time is 30 seconds. The time can be from 1 through 600 seconds. To modify this value, include the **retry-timer** statement:

retry-timer *seconds*;

You can configure this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

By default, no limit is set to the number of times an ingress router attempts to establish or reestablish a connection to the egress router using the primary path. To limit the number of attempts, include the **retry-limit** statement:

retry-limit *number*;

You can configure this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

The limit can be a value up to 10,000. When the retry limit is exceeded, no more attempts are made to establish a path connection. At this point, intervention is required to restart the primary path.

If you set a retry limit, it is reset to 1 each time a successful primary path is created.

Configuring LSP Metrics

The LSP metric is used to indicate the ease or difficulty of sending traffic over a particular LSP. Lower LSP metric values (lower cost) increase the likelihood of an LSP being used. Conversely, high LSP metric values (higher cost) decrease the likelihood of an LSP being used.

The LSP metric can be specified dynamically by the router or explicitly by the user as described in the following sections:

- [Configuring Dynamic LSP Metrics on page 230](#)
- [Configuring Static LSP Metrics on page 231](#)

Configuring Dynamic LSP Metrics

If no specific metric is configured, an LSP attempts to track the IGP metric toward the same destination (the **to** address of the LSP). IGP includes OSPF, IS-IS, Routing Information Protocol (RIP), and static routes. BGP and other RSVP or LDP routes are excluded.

For example, if the OSPF metric toward a router is 20, all LSPs toward that router automatically inherit metric 20. If the OSPF toward a router later changes to a different

value, all LSP metrics change accordingly. If there are no IGP routes toward the router, the LSP raises its metric to 65,535.

Note that in this case, the LSP metric is completely determined by IGP; it bears no relationship to the actual path the LSP is currently traversing. If LSP reroutes (such as through reoptimization), its metric does not change, and thus it remains transparent to users. Dynamic metric is the default behavior; no configuration is required.

Configuring Static LSP Metrics

You can manually assign a fixed metric value to an LSP. Once configured with the **metric** statement, the LSP metric is fixed and cannot change:

metric *number*;

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit protocols mpls **static-label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **static-label-switched-path** *lsp-name*]

The LSP metric has several uses:

- When there are parallel LSPs with the same egress router, the metrics are compared to determine which LSP has the lowest metric value (the lowest cost) and therefore the preferred path to the destination. If the metrics are the same, the traffic is shared.

Adjusting the metric values can force traffic to prefer some LSPs over others, regardless of the underlying IGP metric.

- When an IGP shortcut is enabled (see [“Using Labeled-Switched Paths to Augment SPF to Compute IGP Shortcuts” on page 34](#)), an IGP route might be installed in the routing table with an LSP as the next hop, if the LSP is on the shortest path to the destination. In this case, the LSP metric is added to the other IGP metrics to determine the total path metric. For example, if an LSP whose ingress router is X and egress router is Y is on the shortest path to destination Z, the LSP metric is added to the metric for the IGP route from Y to Z to determine the total cost of the path. If several LSPs are potential next hops, the total metrics of the paths are compared to determine which path is preferred (that is, has the lowest total metric). Or, IGP paths and LSPs leading to the same destination could be compared by means of the metric value to determine which path is preferred.

By adjusting the LSP metric, you can force traffic to prefer LSPs, prefer the IGP path, or share the load among them.

- If router X and Y are BGP peers and if there is an LSP between them, the LSP metric represents the total cost to reach Y from X. If for any reason the LSP reroutes, the underlying path cost might change significantly, but X’s cost to reach Y remains the same (the LSP metric), which allows X to report through a BGP multiple exit

discriminator (MED) a stable metric to downstream neighbors. As long as Y remains reachable through the LSP, no changes are visible to downstream BGP neighbors.

It is possible to configure IS-IS to ignore the configured LSP metric by including the **ignore-lsp-metrics** statement at the **[edit protocols isis traffic-engineering shortcuts]** hierarchy level. This statement removes the mutual dependency between IS-IS and MPLS for path computation. For more information, see the *Junos OS Routing Protocols Library for Routing Devices*.

Configuring CSPF Tie Breaking

When selecting a path for an LSP, CSPF uses a tie-breaking process if there are several equal-cost paths. For information about how CSPF selects a path, see [“How CSPF Selects a Path” on page 31](#).

You can configure one of the following statements (you can only configure one of these statements at a time) to alter the behavior of CSPF tie-breaking:

- By default, a random tie-breaking rule for CSPF is used to select a path from the set of equal-cost paths. However, you can also explicitly configure this behavior using the **random** statement:

random;

- To prefer the path with the least-utilized links, include the **least-fill** statement:

least-fill;

- To prefer the path with the most-utilized links, include the **most-fill** statement:

most-fill;

You can include each of these statements at the following hierarchy levels:

- **[edit protocols mpls label-switched-path *lsp-name*]**
- **[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]**

Related Documentation

- [How CSPF Selects a Path on page 31](#)

Configuring Load Balancing Based on MPLS Labels on DPC I-Chip-Based Hardware

Load balancing occurs on a per-packet basis for MPLS flows on supported platforms. Entropy, or random distribution, is essential for the uniform distribution of packets to their next hops. By default, when load balancing is used to help distribute traffic, Junos OS employs a hash algorithm to select a next-hop address to install into the forwarding table. Whenever the set of next hops for a destination changes, the next-hop address is reselected by means of the hash algorithm. You can configure how the hash algorithm is used to load-balance traffic across a set of equal-cost label switched paths (LSPs).

To ensure entropy for VPLS & VPWS traffic, Junos OS can create a hash based on data from the IP header and as many as three MPLS labels (the so-called top labels).

In some cases, as the number of network feature that use labels grows (such as MPLS Fast Reroute, and RFC 3107, RSVP and VPN) data in the top three labels can become static and thus not a sufficient source for entropy. Load balancing can become skewed as a result, or the incidence of out-of-order packet delivery may rise. For these cases, labels from the bottom of the label stack can be used (see Table 1, below for qualifications). Top labels and bottom labels cannot be used at the same time.



NOTE: MPC cards do not support the regular hash key configuration. For the MPC-based hash key configuration to be effective, you need an **enhanced-hash-key** configuration.

Load balancing is used to evenly distribute traffic when the following conditions apply:

- There are multiple equal-cost next hops over different interfaces to the same destination.
- There is a single next hop over an aggregated interface.

An LSP tends to load-balance its placement by randomly selecting one of the equal-cost next hops and using it exclusively. The random selection is made independently at each transit router, which compares Interior Gateway Protocol (IGP) metrics alone. No consideration is given to bandwidth or congestion levels.

This feature applies to aggregated Ethernet and aggregated SONET/SDH interfaces as well as multiple equal-cost MPLS next hops. In addition, on the T Series, MX Series, M120, and M320 routers only, you can configure load balancing for IPv4 traffic over Layer 2 Ethernet pseudowires. You can also configure load balancing for Ethernet pseudowires based on IP information. The option to include IP information in the hash key provides support for Ethernet circuit cross-connect (CCC) connections.

To load-balance based on the MPLS label information, configure the **family mpls** statement:

```
[edit forwarding-options hash-key]
family mpls {
  all-labels;
  bottom-label-1;
  bottom-label-2;
  bottom-label-3;
  label-1;
  label-2;
  label-3;
  no-labels;
  no-label-1-exp;
  payload {
    ether-pseudowire;
    ip {
      disable;
      layer-3-only;
      port-data {
        destination-lsb;
        destination-msb;
      }
    }
  }
}
```

```

        source-lsb;
        source-msb;
    }
}
}

```

You can include this statement at the following hierarchy levels:

- **[edit forwarding-options hash-key]**

[Table 3 on page 234](#) provides detailed information about all of the possible MPLS LSP load-balancing options.

Table 3: MPLS LSP Load Balancing Options

Statement	Supported Platforms	MPLS LSP Load Balancing Options
all-labels	PTX Series	Up to eight MPLS labels are included in the hash key to identify the uniqueness of a flow in the Packet Forwarding Engine. This value is set by default.
bottom-label-1	MX Series with DPC (I-Chip). Not supported on M10i, M7i, and M120.	Uses the bottom-most label for calculating the hash key, for example if the top labels do not provide sufficient variable for the required level of entropy.
bottom-label-2	MX Series with DPC (I-Chip). Not supported on M10i, M7i, and M120.	Uses the second label from the bottom for calculating the hash key, for example if the top labels do not provide sufficient variable for the required level of entropy.
bottom-label-3	MX Series with DPC (I-Chip). Not supported on M10i, M7i, and M120.	Uses the third label from the bottom for calculating the hash key, for example if the top labels do not provide sufficient variable for the required level of entropy.
label-1	M Series, MX Series, T Series	Include the first label in the hash key. Use this option for single label packets.
label-2	M Series, MX Series, T Series	Include the second label in the hash key. You must also configure the label-1 option. The entire first label and the first 16 bits of the second label are used in the hash key.
label-3	M Series, MX Series, T Series	Include the third label in the hash key. You must also configure the label-1 option and the label-2 option.
no-labels	All	Excludes MPLS labels from the hash key.
no-label-1-exp	M Series, MX Series, T Series	Excludes the EXP bit of the top label from the hash key. You must also configure the label-1 option. For Layer 2 VPNs, the router could encounter a packet reordering problem. When a burst of traffic pushes the customer traffic bandwidth to exceed its limits, the traffic might be affected in mid flow. Packets might be reordered as a result. By excluding the EXP bit from the hash key, you can avoid this reordering problem.

Table 3: MPLS LSP Load Balancing Options (*continued*)

Statement	Supported Platforms	MPLS LSP Load Balancing Options
payload	All	Allows you to configure which parts of the IP packet payload to include in the hash key. For the PTX Series Packet Transport Router, this value is set by default.
disable	PTX Series	Exclude IP payload from the hash key.
ether-pseudowire	M120, M320, MX Series, T Series	Load-balance IPv4 traffic over Layer 2 Ethernet pseudowires.
ip	All	Include the IPv4 or IPv6 address in the hash key. You must also configure either label-l or no-labels .
layer-3-only	All	Include only the Layer 3 IP information in the hash key. Excludes all of the port-data bytes from the hash key.
port-data	M Series, MX Series, T Series	Include the source and destination port field information. By default, the most significant byte and least significant byte of the source and destination port fields are used in the hash key. To select specific bytes to use in the hash key, include one or more of the source-msb , source-lsb , destination-msb , and destination-lsb options at the [edit forwarding-options hash-key family mpls payload ip port-data] hierarchy level. To prevent all four bytes from being hashed, include the layer-3-only statement at the [edit forwarding-options hash-key family mpls payload ip] hierarchy level.
destination-lsb	M Series, MX Series, T Series	Include the least significant byte of the destination port in the hash key. Can be combined with any of the other port-data options.
destination-msb	M Series, MX Series, T Series	Include the most significant byte of the destination port in the hash key. Can be combined with any of the other port-data options.
source-lsb	M Series, MX Series, T Series	Include the least significant byte of the source port in the hash key. Can be combined with any of the other port-data options.
source-msb	M Series, MX Series, T Series	Include the most significant byte of the source port in the hash key. Can be combined with any of the other port-data options.

The following examples illustrate ways in which you can configure MPLS LSP load balancing:

- To include the IP address as well as the first label in the hash key:
 - For M Series, MX Series, and T Series routers, configure the **label-l** statement and the **ip** option for the **payload** statement at the **[edit forwarding-options hash-key family mpls]** hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-l;
payload {
  ip;
}
```

- For PTX Series Packet Transport Routers, the **all-labels** and **ip payload** options are configured by default, so no configuration is necessary.
- (M320 and T Series routers only) To include the IP address as well as both the first and second labels in the hash key, configure the **label-1** and **label-2** options and the **ip** option for the **payload** statement at the **[edit forwarding-options hash-key family mpls]** hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
label-2;
payload {
  ip;
}
```



NOTE: You can include this combination of statements on M320 and T Series routers only. If you include them on an M Series Multiservice Edge Router, only the first MPLS label and the IP payload are used in the hash key.

- For T Series routers, ensure proper load balancing by including the **label-1**, **label-2**, and **label-3** options at the **[edit forwarding-options hash-key family mpls]** hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
label-2;
label-3;
```

- (M Series, MX Series, and T Series routers only) For Layer 2 VPNs, the router could encounter a packet reordering problem. When a burst of traffic pushes the customer traffic bandwidth to exceed its limits, the traffic might be affected in mid flow. Packets might be reordered as a result. By excluding the EXP bit from the hash key, you can avoid this reordering problem. To exclude the EXP bit of the first label from the hash calculations, include the **no-label-1-exp** statement at the **[edit forwarding-options hash-key family mpls]** hierarchy level:

```
[edit forwarding-options hash-key family mpls]
label-1;
no-label-1-exp;
payload {
  ip;
}
```

Related Documentation

- *Configuring Load Balancing for Ethernet Pseudowires*

Disabling Normal TTL Decrementing

By default, the time-to-live (TTL) field value in the packet header is decremented by 1 for every hop the packet traverses in the LSP, thereby preventing loops. If the TTL field value reaches 0, packets are dropped, and an Internet Control Message Protocol (ICMP) error packet is sent to the originating router.

If the normal TTL decrement is disabled, the TTL field of IP packets entering LSPs are decremented by only 1 on transiting the LSP, making the LSP appear as a one-hop router to diagnostic tools, such as **traceroute**. Decrementing the TTL field by 1 is done by the ingress router, which pushes a label on IP packets with the TTL field in the label initialized to 255. The label's TTL field value is decremented by 1 for every hop the MPLS packet traverses in the LSP. On the penultimate hop of the LSP, the router pops the label but does not write the label's TTL field value to the IP packet's TTL field. Instead, when the IP packet reaches the egress router, the IP packet's TTL field value is decremented by 1.

When you use **traceroute** to diagnose problems with an LSP from outside that LSP, **traceroute** sees the ingress router, even though the egress router performs the TTL decrement. The behavior of **traceroute** is different if it is initiated from the ingress router of the LSP. In this case, the egress router would be the first router to respond to **traceroute**.

You can disable normal TTL decrementing in an LSP so that the TTL field value does not reach 0 before the packet reaches its destination, thus preventing the packet from being dropped. You can also disable normal TTL decrementing to make the MPLS cloud appear as a single hop, thereby hiding the network topology.

There are two ways to disable TTL decrementing:

- On the ingress of the LSP, if you include the **no-decrement-ttl** statement, the ingress router negotiates with all downstream routers using a proprietary RSVP object, to ensure all routers are in agreement. If negotiation succeeds, the whole LSP behaves as one hop to transit IP traffic.

no-decrement-ttl;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.



NOTE: The RSVP object is proprietary to the Junos OS and might not work with other software. This potential incompatibility applies only to RSVP-signaled LSPs. When you include the **no-decrement-ttl** statement, TTL hiding can be enforced on a per-LSP basis.

- On the ingress router, you can include the **no-propagate-ttl** statement. The **no-propagate-ttl** statement applies to all LSPs, regardless of whether they are RSVP-signaled or LDP-signaled. Once set, all future LSPs traversing through this router behave as a single hop to IP packets. LSPs established before you configure this statement are not affected.

no-propagate-ttl;

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

The operation of the **no-propagate-ttl** statement is interoperable with other vendors' equipment. However, you must ensure that all routers are configured identically.

To configure the TTL behavior for a single VRF routing instance, include the **no-vrf-propagate-ttl** or the **vrf-propagate-ttl** statement in the routing instance configuration at the `[edit routing-instances instance-name]` hierarchy level. The **no-vrf-propagate-ttl** or the **vrf-propagate-ttl** statement overrides the behavior configured globally for the router. If the router is operating in default mode with normal TTL decrementing, the **no-vrf-propagate-ttl** overrides the global behavior for the routing instance on which the **no-vrf-propagate-ttl** statement is configured.

Related Documentation

- *Example: Disabling Normal TTL Decrementing in a VRF Routing Instance (on Layer 3 VPNs Feature Guide for Routing Devices in the Junos VPNs Configuration Guide*

Configuring MPLS Soft Preemption

Soft preemption attempts to establish a new path for a preempted LSP before tearing down the original LSP. The default behavior is to tear down a preempted LSP first, signal a new path, and then reestablish the LSP over the new path. In the interval between when the path is taken down and the new LSP is established, any traffic attempting to use the LSP is lost. Soft preemption prevents this type of traffic loss. The trade-off is that during the time when an LSP is being soft preempted, two LSPs with their corresponding bandwidth requirements are used until the original path is torn down.

MPLS soft preemption is useful for network maintenance. For example, you can move all LSPs away from a particular interface, then take the interface down for maintenance without interrupting traffic. MPLS soft preemption is described in detail in RFC 5712, *MPLS Traffic Engineering Soft Preemption*.

Soft preemption is a property of the LSP and is disabled by default. You configure it at the ingress of an LSP by including the **soft-preemption** statement:

```
soft-preemption;
```

You can include this statement at the following hierarchy levels:

- `[edit protocols mpls label-switched-path lsp-name]`
- `[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name]`

You can also configure a timer for soft preemption. The timer designates the length of time the router should wait before initiating a hard preemption of the LSP. At the end of the time specified, the LSP is torn down and resignaled. The soft-preemption cleanup timer has a default value of 30 seconds; the range of permissible values is 0 through 180 seconds. A value of 0 means that soft preemption is disabled. The soft-preemption cleanup timer is global for all LSPs.

Configure the timer by including the **cleanup-timer** statement:

```
cleanup-timer seconds;
```

You can include this statement at the following hierarchy levels:

- `[edit protocols rsvp preemption soft-preemption]`

- [edit logical-systems *logical-system-name* protocols RSVP **preemption soft-preemption**]



NOTE: Soft preemption cannot be configured on LSPs for which secondary paths or fast reroute has been configured. The configuration fails to commit. However, you can enable soft preemption in conjunction with node and link protection.

Disabling Constrained-Path LSP Computation

If the IGP is a link-state protocol (such as IS-IS or OSPF) and supports extensions that allow the current bandwidth reservation on each router's link to be reported, constrained-path LSPs are computed by default.

The Junos implementations of IS-IS and OSPF include the extensions that support constrained-path LSP computation.

- IS-IS—These extensions are enabled by default. To disable this support, include the **disable** statement at the [edit protocols isis traffic-engineering] hierarchy level, as discussed in the *Junos OS Routing Protocols Library for Routing Devices*.
- OSPF—These extensions are disabled by default. To enable this support, include the **traffic-engineering** statement in the configurations of all routers running OSPF, as described in the *Junos OS Routing Protocols Library for Routing Devices*.

If IS-IS is enabled on a router or you enable OSPF traffic engineering extensions, MPLS performs the constrained-path LSP computation by default. For information about how constrained-path LSP computation works, see [“Constrained-Path LSP Computation” on page 29](#).

Constrained-path LSPs have a greater chance of being established quickly and successfully for the following reasons:

- The LSP computation takes into account the current bandwidth reservation.
- Constrained-path LSPs reroute themselves away from node failures and congestion.

When constrained-path LSP computation is enabled, you can configure the LSP so that it is periodically reoptimized, as described in [“Optimizing Signaled LSPs” on page 251](#).

When an LSP is being established or when an existing LSP fails, the constrained-path LSP computation is repeated periodically at the interval specified by the retry timer until the LSP is set up successfully. Once the LSP is set up, no recomputation is done. For more information about the retry timer, see [“Configuring the Connection Between Ingress and Egress Routers” on page 229](#).

By default, constrained-path LSP computation is enabled. You might want to disable constrained-path LSP computation when all nodes do not support the necessary traffic engineering extensions. To disable constrained-path LSP computation, include the **no-cspf** statement:

no-cspf;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

If you disable constrained-path LSP computation on LSPs by configuring the **no-cspf** statement and then attempt to advertise other LSPs with lower metrics than the IGP from this router in either IS-IS or OSPF, new LSPs cannot be established.

Configuring Administrative Groups for LSPs

Administrative groups, also known as link coloring or resource class, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. You can use administrative groups to implement a variety of policy-based LSP setups.

Administrative groups are meaningful only when constrained-path LSP computation is enabled.

You can assign up to 32 names and values (in the range 0 through 31), which define a series of names and their corresponding values. The administrative names and values must be identical across all routers within a single domain.



NOTE: The administrative value is distinct from the priority. You configure the priority for an LSP using the **priority** statement. See “[Configuring Priority and Preemption for LSPs](#)” on page 250.

To configure administrative groups, follow these steps:

1. Define multiple levels of service quality by including the **admin-groups** statement:

```
admin-groups {  
  group-name group-value;  
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

The following configuration example illustrates how you might configure a set of administrative names and values for a domain:

```
[edit protocols mpls]  
admin-groups {  
  gold 1;  
  silver 2;  
  copper 3;  
  best-effort 4;  
}
```

2. Define the administrative groups to which an interface belongs. You can assign multiple groups to an interface. Include the **interface** statement:

```
interface interface-name {
  admin-group [ group-names ];
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

If you do not include the **admin-group** statement, an interface does not belong to any group.

IGPs use the group information to build link-state packets, which are then flooded throughout the network, providing information to all nodes in the network. At any router, the IGP topology, as well as administrative groups of all the links, is available.

Changing the interface's administrative group affects only new LSPs. Existing LSPs on the interface are not preempted or recomputed to keep the network stable. If LSPs need to be removed because of a group change, issue the **clear RSVP session** command.

3. Configure an administrative group constraint for each LSP or for each primary or secondary LSP path. Include the **label-switched-path** statement:

```
label-switched-path lsp-name {
  to address;
  ...
  admin-group {
    exclude [ group-names ];
    include-all [ group-names ];
    include-any [ group-names ];
  }
  primary path-name {
    admin-group {
      exclude [ group-names ];
      include-all [ group-names ];
      include-any [ group-names ];
    }
  }
  secondary path-name {
    admin-group {
      exclude [ group-names ];
      include-all [ group-names ];
      include-any [ group-names ];
    }
  }
}
```

You can include the **label-switched-path** statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

If you omit the **include-all**, **include-any**, or **exclude** statements, the path computation proceeds unchanged. The path computation is based on the constrained-path LSP computation. For information about how the constrained-path LSP computation is calculated, see [“How CSPF Selects a Path” on page 31](#).



NOTE: Changing the LSP's administrative group causes an immediate recomputation of the route; therefore, the LSP might be rerouted.

Configuring Extended Administrative Groups for LSPs

In MPLS traffic engineering, a link can be configured with a set of administrative groups (also known as colors or resource classes). Administrative groups are carried in the interior gateway protocol (IGP) (OSPFv2 and IS-IS) as a 32-bit value assigned to each link. Juniper Networks routers normally interpret this 32-bit value as a bit mask with each bit representing a group, limiting each network to a total of 32 distinct administrative groups (value range 0 through 31).

You configure extended administrative groups, represented by a 32-bit value, expanding the number of administrative groups supported in the network beyond just 32. The original range of values available for administrative groups is still supported for backwards compatibility.

The extended administrative groups configuration accepts a set of interfaces with a corresponding set of extended administrative group names. It converts the names into a set of 32-bit values and propagates this information into the IGP. The extended administrative group values are global and must be identically configured on all the supported routers participating in the network. The domain-wide extended administrative groups database, learned from other routers through IGP flooding, is used by Constrained Shortest Path First (CSPF) for path computation.

The following procedure describes how to configure extended administrative groups:

1. Configure the `admin-groups-extended-range` statement:

```
admin-groups-extended-range {  
    maximum maximum-number;  
    minimum minimum-number;  
}
```

You can include this statement at the following hierarchy levels:

- `[edit routing-options]`
- `[edit logical-systems logical-system-name routing-options]`

The `admin-groups-extended-range` statement includes the `minimum` and `maximum` options. The range maximum must be greater than the range minimum.

2. Configure the `admin-groups-extended` statement:

```
admin-groups-extended group-name {  
    group-value group-identifier;  
}
```

You can include this statement at the following hierarchy levels:

- `[edit routing-options]`

- [edit logical-systems *logical-system-name* routing-options]

The **admin-groups-extended** statement enables you to configure a group name and group value for the administrative group. The group value must be within the range of values configured using the **admin-groups-extended-range** statement.

3. The extended administrative groups for an MPLS interface consist of the set of extended administrative group names assigned for the interface. The interface extended administrative group names must be configured for the global extended administrative groups.

To configure an extended administrative group for an MPLS interface, specify the administrative group name within the MPLS interface configuration using the **admin-groups-extended** statement:

```
admin-groups-extended group-name;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls interface *interface-name*]
- [edit logical-systems *logical-system-name* protocols mpls interface *interface-name*]

4. The LSP extended administrative groups define the set of include and exclude constraints for an LSP and for a path's primary and secondary paths. The extended administrative group names must be configured for the global extended administrative groups.

To configure extended administrative groups for an LSP, include the **admin-group-extended** statement at an LSP hierarchy level:

```
admin-group-extended {
  apply-groups group-value;
  apply-groups-except group-value;
  exclude group-value;
  include-all group-value;
  include-any group-value;
}
```

The **admin-group-extended** statement includes the following options: **apply-groups**, **apply-groups-except**, **exclude**, **include-all**, and **include-any**. Each option enables you to configure one or more extended administrative groups.

For the list of the hierarchy levels at which you can configure this statement, see the statement summary for this statement.

5. To display the currently configured extended administrative groups, issue the **show mpls admin-groups-extended** command.

Related Documentation

- [Configuring Administrative Groups for LSPs on page 240](#)

Configuring Preference Values for LSPs

As an option, you can configure multiple LSPs between the same pair of ingress and egress routers. This is useful for balancing the load among the LSPs because all LSPs,

by default, have the same preference level. To prefer one LSP over another, set different preference levels for individual LSPs. The LSP with the lowest preference value is used. The default preference for RSVP LSPs is 7 and for LDP LSPs is 9. These preference values are lower (more preferred) than all learned routes except direct interface routes.

To change the default preference value, include the **preference** statement:

```
preference preference;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Disabling Path Route Recording by LSPs

The Junos implementation of RSVP supports the Record Route object, which allows an LSP to actively record the routers through which it transits. You can use this information for troubleshooting and to prevent routing loops. By default, path route information is recorded. To disable recording, include the **no-record** statement:

```
no-record;
```

For a list of hierarchy levels at which you can include the **record** and **no-record** statements, see the statement summary section for the statement.

Configuring Class of Service for MPLS LSPs

The following sections provide an overview of MPLS class of service (CoS) and describe how to configure the MPLS CoS value:

- [Class of Service for MPLS Overview on page 244](#)
- [Configuring the MPLS CoS Bits on page 245](#)
- [Rewriting IEEE 802.1p Packet Headers with the MPLS CoS Value on page 246](#)

Class of Service for MPLS Overview

When IP traffic enters an LSP tunnel, the ingress router marks all packets with a CoS value, which is used to place the traffic into a transmission priority queue. On the router, for SDH/SONET and T3 interfaces, each interface has four transmit queues. The CoS value is encoded as part of the MPLS header and remains in the packets until the MPLS header is removed when the packets exit from the egress router. The routers within the LSP utilize the CoS value set at the ingress router. The CoS value is encoded by means of the CoS bits (also known as the EXP or experimental bits). For more information, see [“MPLS Label Allocation” on page 26](#).

MPLS class of service works in conjunction with the router’s general CoS functionality. If you do not configure any CoS features, the default general CoS settings are used. For MPLS class of service, you might want to prioritize how the transmit queues are serviced by configuring weighted round-robin, and to configure congestion avoidance using random early detection (RED). The general CoS features are described in the *Class of Service Feature Guide for Routing Devices*.

Configuring the MPLS CoS Bits

When traffic enters an LSP tunnel, the CoS bits in the MPLS header are set in one of two ways:

- The number of the output queue into which the packet was buffered and the packet loss priority (PLP) bit are written into the MPLS header and are used as the packet's CoS value. This behavior is the default, and no configuration is required. The *Class of Service Feature Guide for Routing Devices* explains the IP CoS values, and summarizes how the CoS bits are treated.
- You set a fixed CoS value on all packets entering the LSP tunnel. A fixed CoS value means that all packets entering the LSP receive the same class of service.

To set a fixed CoS value on all packets entering the LSP, include the **class-of-service** statement:

```
class-of-service cos-value;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The CoS value set using the **class-of-service** statement at the **[edit protocols mpls]** hierarchy level supersedes the CoS value set at the **[edit class-of-service]** hierarchy level for an interface. Effectively, the CoS value configured for an LSP overrides the CoS value set for an interface.

The CoS value can be a decimal number from 0 through 7. This number corresponds to a 3-bit binary number. The high-order 2 bits of the CoS value select which transmit queue to use on the outbound interface card.

The low-order bit of the CoS value is treated as the PLP bit and is used to select the RED drop profile to use on the output queue. If the low-order bit is 0, the non-PLP drop profile is used, and if the low-order bit is 1, the PLP drop profile is used. It is generally expected that RED will more aggressively drop packets that have the PLP bit set. For more information about RED and drop profiles, see the *Class of Service Feature Guide for Routing Devices*.



NOTE: Configuring the PLP drop profile to drop packets more aggressively (for example, setting the CoS value from 6 to 7) decreases the likelihood of traffic getting through.

Table 4 on page 246 summarizes how MPLS CoS values correspond to the transmit queue and PLP bit. Note that in MPLS, the mapping between the CoS bit value and the output queue is hard-coded. You cannot configure the mapping for MPLS; you can configure it only for IPv4 traffic flows, as described in the *Class of Service Feature Guide for Routing Devices*.

Table 4: MPLS CoS Values

MPLS CoS Value	Bits	Transmit Queue	PLP Bit
0	000	0	Not set
1	001	0	Set
2	010	1	Not set
3	011	1	Set
4	100	2	Not set
5	101	2	Set
6	110	3	Not set
7	111	3	Set

Because the CoS value is part of the MPLS header, the value is associated with the packets only as they travel through the LSP tunnel. The value is not copied back to the IP header when the packets exit from the LSP tunnel.

Rewriting IEEE 802.1p Packet Headers with the MPLS CoS Value

For Ethernet interfaces installed on a T Series router or an M320 router with a peer connection to an M Series router or a T Series router, you can rewrite both MPLS CoS and IEEE 802.1p bits to a configured value (the MPLS CoS bits are also known as the EXP or experimental bits). Rewriting these bits allows you to pass the configured value to the Layer 2 VLAN path. To rewrite both the MPLS CoS and IEEE 802.1p bits, you must include the EXP and IEEE 802.1p rewrite rules in the class of service interface configuration. The EXP rewrite table is applied when you configure the IEEE 802.1p and EXP rewrite rules.

For information about how to configure the EXP and IEEE 802.1p rewrite rules, see the *Class of Service Feature Guide for Routing Devices*.

Achieving a Make-Before-Break, Hitless Switchover for LSPs

Adaptive label-switched paths (LSPs) might need to establish a new LSP instance and transfer traffic from an old LSP instance onto the new LSP instance before tearing down the old one. This type of configuration is referred to as a *make before break* (MBB).

RSVP-TE is a protocol used to establish LSPs in MPLS networks. The Junos OS implementation of RSVP-TE to achieve a hitless (no traffic loss) MBB switchover has relied on configuring the timer values in the following configuration statements:

- **optimize-switchover-delay**—Amount of time to wait before switching to the new LSP instance.
- **optimize-hold-dead-delay**—Amount of time to wait after switchover and before deletion of the old LSP instance.

Both the **optimize-switchover-delay** and **optimize-hold-dead-delay** statements apply to all LSPs that use the make-before-break behavior for LSP setup and teardown, not just for LSPs for which the **optimize-timer** statement has also been configured. The following MPLS features cause LSPs to be set up and torn down using make-before-break behavior:

- Adaptive LSPs
- Automatic bandwidth allocation
- BFD for LSPs
- Graceful Routing Engine switchover
- Link and node protection
- Nonstop active routing
- Optimized LSPs
- Point-to-multipoint (P2MP) LSPs
- Soft preemption
- Standby secondary paths

Both the **optimize-switchover-delay** and **optimize-hold-dead-delay** statements when configured add an artificial delay to the MBB process. The value of the **optimize-switchover-delay** statement varies with the size of the Explicit Route Objects (EROs). An ERO is an extension to RSVP that allows an RSVP PATH message to traverse an explicit sequence of routers that is independent of conventional shortest-path IP routing. The value of the **optimize-switchover-delay** statement also depends on the CPU load on each of the routers on the path. Customers set the **optimize-switchover-delay** statement by trial and error.

The value of the **optimize-hold-dead-delay** statement depends on how fast the ingress router moves all application prefixes to point to the new LSP. This is determined by the Packet Forwarding Engine load, which can vary from platform to platform. Customers have to set the **optimize-hold-dead-delay** statement by trial and error.

However, as of Release 15.1, Junos OS is able to achieve a hitless MBB switchover without configuring the artificial delays that such timer values introduce.

This topic summarizes the three methods of achieving a MBB switchover from an old LSP to a new LSP using Junos OS:

- [Specifying the Amount of Time the Router Waits to Switch Over to New Paths on page 247](#)
- [Specifying the Amount of Time to Delay the Tear Down of Old Paths on page 248](#)
- [Achieving a Hitless, MBB Switchover Without Artificial Delays on page 248](#)

Specifying the Amount of Time the Router Waits to Switch Over to New Paths

To specify the amount of time the router waits to switch over LSP instances to newly optimized paths, use the **optimize-switchover-delay** statement. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need

to configure this statement on transit or egress routers). The timer in this statement helps to ensure that the new optimized paths have been established before traffic is switched over from the old paths. This timer can only be enabled or disabled for all of the LSPs configured on the router.

To configure the amount of time the router waits to switch over LSP instances to newly optimized paths, specify the time in seconds by using the **optimize-switchover-delay** statement:

```
optimize-switchover-delay seconds;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

Specifying the Amount of Time to Delay the Tear Down of Old Paths

To specify the amount of time to delay the tear down of old paths after the router has switched traffic to new optimized paths, use the **optimize-hold-dead-delay** statement. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The timer in this statement helps to ensure that old paths are not torn down before all routes have been switched over to the new optimized paths. This timer can be enabled for specific LSPs or for all of the LSPs configured on the router.

To configure the amount of time in seconds to delay the tear down of old paths after the router has switched traffic to new optimized paths, use the **optimize-hold-dead-delay** statement:

```
optimize-hold-dead-delay seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Achieving a Hitless, MBB Switchover Without Artificial Delays

As of Junos OS Release 15.1, there is another way to relinquish the old LSP instances after MBB switchover without relying on the arbitrary time intervals set up by the **optimize-switchover-delay** or **optimize-hold-dead-delay** statement. For example, if you use the **optimize-hold-dead-delay** statement, you configure a time you think it is safe to wait before tearing down the old LSP instance after MBB. However, some routes might still be in the process of shifting to the new instance. Tearing down the old LSP instance prematurely results in one of the transit nodes dropping the traffic for those routes that have not shifted to the new LSP instance.

To avoid traffic loss, instead of using the **optimize-switchover-delay** statement, you can use MPLS-OAM (lsp ping), which confirms that the LSP data plane is established end-to-end. Instead of using the **optimize-hold-dead-delay** statement, you can use a feedback mechanism from the rpd infrastructure that confirms that all prefixes referring to the old LSP have been switched over. The feedback mechanism is sourced from the Tag library and relies on the routing protocol process (rpd) infrastructure to determine

when all the routes using the old LSP instance have fully shifted to the new LSP instance after MBB switchover.

The feedback mechanism is always in place, and it is optional. Configure the **optimize-adaptive-teardown** statement to have the feedback mechanism used during MBB switchover. This feature is not supported for RSVP point-to-multipoint (P2MP) LSP instances. Global configuration of the **optimize-adaptive-teardown** statement only affects the point-to-point LSPs that are configured in the system.

You only need to configure the **optimize-adaptive-teardown** statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). This feedback mechanism ensures that old paths are not torn down before all routes have been switched over to the new optimized paths. The global configuration of this configuration statement affects only the point-to-point LSPs that are configured in the system.

```
optimize-adaptive-teardown {
  p2p:
}
```

You can include this statement at the **[edit protocols mpls]** hierarchy level.

Related Documentation

- [Configuring Adaptive LSPs on page 249](#)
- [Configuring Automatic Bandwidth Allocation for LSPs on page 257](#)
- [Configuring MPLS Soft Preemption on page 238](#)
- [Configuring the Smart Optimize Timer for LSPs on page 255](#)
- [Configuring Hot Standby of Secondary Paths for LSPs on page 267](#)

Configuring Adaptive LSPs

An LSP occasionally might need to reroute itself for these reasons:

- The continuous reoptimization process is configured with the **optimize-timer** statement.
- The current path has connectivity problems.
- The LSP is preempted by another LSP configured with the **priority** statement and is forced to reroute.
- The explicit-path information for an active LSP is modified, or the LSP's bandwidth is increased.

You can configure an LSP to be *adaptive* when it is attempting to reroute itself. When it is adaptive, the LSP holds onto existing resources until the new path is successfully established and traffic has been cut over to the new LSP. To retain its resources, an adaptive LSP does the following:

- Maintains existing paths and allocated bandwidths—This ensures that the existing path is not torn down prematurely and allows the current traffic to continue flowing while the new path is being set up.

- Avoids double-counting for links that share the new and old paths—Double-counting occurs when an intermediate router does not recognize that the new and old paths belong to the same LSP and counts them as two separate LSPs, requiring separate bandwidth allocations. If some links are close to saturation, double-counting might cause the setup of the new path to fail.

By default, adaptive behavior is disabled. You can include the **adaptive** statement in two different hierarchy levels.

If you specify the **adaptive** statement at the LSP hierarchy levels, the adaptive behavior is enabled on all primary/secondary paths of the LSP. This means both the primary and secondary paths share the same bandwidth on common links.

To configure adaptive behavior for all LSP paths, include the **adaptive** statement in the LSP configuration:

```
adaptive;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

If you specify the **adaptive** statement at the [edit protocols mpls **label-switched-path** *lsp-name* (**primary** | **secondary**) *path-name*] hierarchy level, adaptive behavior is enabled only on the path on which it is specified. Bandwidth double-counting occurs between different paths. However, if you also have the **adaptive** statement configured at the [edit protocols mpls **label-switched-path** *lsp-name*] hierarchy level, it overrides the adaptive behavior of each individual path.

To configure adaptive behavior for either the primary or secondary level, include the **adaptive** statement:

```
adaptive;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name* (**primary** | **secondary**) *path-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name* (**primary** | **secondary**) *path-name*]

Configuring Priority and Preemption for LSPs

When there is insufficient bandwidth to establish a more important LSP, you might want to tear down a less important existing LSP to free the bandwidth. You do this by preempting the existing LSP.

Whether an LSP can be preempted is determined by two properties associated with the LSP:

- Setup priority—Determines whether a new LSP that preempts an existing LSP can be established. For preemption to occur, the setup priority of the new LSP must be higher

than that of the existing LSP. Also, the act of preempting the existing LSP must produce sufficient bandwidth to support the new LSP. That is, preemption occurs only if the new LSP can be set up successfully.

- **Reservation priority**—Determines the degree to which an LSP holds on to its session reservation after the LSP has been set up successfully. When the reservation priority is high, the existing LSP is less likely to give up its reservation, and hence it is unlikely that the LSP can be preempted.

You cannot configure an LSP with a high setup priority and a low reservation priority, because permanent preemption loops might result if two LSPs are allowed to preempt each other. You must configure the reservation priority to be higher than or equal to the setup priority.

The setup priority also defines the relative importance of LSPs on the same ingress router. When the software starts, when a new LSP is established, or during fault recovery, the setup priority determines the order in which LSPs are serviced. Higher-priority LSPs tend to be established first and hence enjoy more optimal path selection.

To configure the LSP's preemption properties, include the **priority** statement:

```
priority setup-priority reservation-priority;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Both **setup-priority** and **reservation-priority** can be a value from 0 through 7. The value 0 corresponds to the highest priority, and the value 7 to the lowest. By default, an LSP has a setup priority of 7 (that is, it cannot preempt any other LSPs) and a reservation priority of 0 (that is, other LSPs cannot preempt it). These defaults are such that preemption does not happen. When you are configuring these values, the setup priority should always be less than or equal to the hold priority.

Optimizing Signaled LSPs

Once an LSP has been established, topology or resources changes might, over time, make the path suboptimal. A new path might have become available that is less congested, has a lower metric, and traverses fewer hops. You can configure the router to recompute paths periodically to determine whether a more optimal path has become available.

If reoptimization is enabled, an LSP can be rerouted through different paths by constrained-path recomputations. However, if reoptimization is disabled, the LSP has a fixed path and cannot take advantage of newly available network resources. The LSP is fixed until the next topology change breaks the LSP and forces a recomputation.

Reoptimization is not related to failover. A new path is always computed when topology failures occur that disrupt an established path.

Because of the potential system overhead involved, you need to carefully control the frequency of reoptimization. Network stability might suffer when reoptimization is enabled. By default, the **optimize-timer** statement is set to 0 (that is, it is disabled).

LSP optimization is meaningful only when constrained-path LSP computation is enabled, which is the default behavior. For more information about constrained-path LSP computation, see [“Disabling Constrained-Path LSP Computation” on page 239](#). Also, LSP optimization is only applicable to ingress LSPs, so it is only necessary to configure the **optimize-timer** statement on the ingress router. The transit and egress routers require no specific configuration to support LSP optimization (other than to have MPLS enabled).

To enable path reoptimization, include the **optimize-timer** statement:

optimize-timer *seconds*;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Once you have configured the **optimize-timer** statement, the reoptimization timer continues its countdown to the configured value even if you delete the **optimize-timer** statement from the configuration. The next optimization uses the new value. You can force the Junos OS to use a new value immediately by deleting the old value, committing the configuration, configuring the new value for the **optimize-timer** statement, and then committing the configuration again.

After reoptimization is run, the result is accepted only if it meets the following criteria:

1. The new path is not higher in IGP metric. (The metric for the old path is updated during computation, so if a recent link metric changed somewhere along the old path, it is accounted for.)
2. If the new path has the same IGP metric, it is not more hops away.
3. The new path does not cause preemption. (This is to reduce the ripple effect of preemption causing more preemption.)
4. The new path does not worsen congestion overall.

The relative congestion of the new path is determined as follows:


- a. The percentage of available bandwidth on each link traversed by the new path is compared to that for the old path, starting from the most congested links.
- b. For each current (old) path, the software stores the four smallest values for bandwidth availability for the links traversed in ascending order.
- c. The software also stores the four smallest bandwidth availability values for the new path, corresponding to the links traversed in ascending order.
- d. If any of the four new available bandwidth values are smaller than any of the corresponding old bandwidth availability values, the new path has at least one link that is more congested than the link used by the old path. Because using the link would cause more congestion, traffic is not switched to this new path.
- e. If none of the four new available bandwidth values is smaller than the corresponding old bandwidth availability values, the new path is less congested than the old path.

When all the above conditions are met, then:

5. If the new path has a lower IGP metric, it is accepted.
6. If the new path has an equal IGP metric and lower hop count, it is accepted.
7. If you choose **least-fill** as a load balancing algorithm, LSPs are load balanced as follows:
 - a. The LSP is moved to a new path that is utilized at least 10% less than the current path. This might reduce congestion on the current path by only a small amount. For example, if an LSP with 1 MB of bandwidth is moved off a path carrying a minimum of 200 MB, congestion on the original path is reduced by less than 1%.
 - b. For **random** or **most-fill** algorithms, this rule does not apply.

The following example illustrates how the **least-fill** load balancing algorithm works.

Figure 30: least-fill Load Balancing Algorithm Example



As shown in [Figure 30 on page 253](#), there are two potential paths for an LSP to traverse from router A to router H, the odd links from L1 through L13 and the even links from L2 through L14. Currently, the router is using the even links as the active path for the LSP. Each link between the same two routers (for example, router A and router B) has the same bandwidth:

- L1, L2 = 10GE
- L3, L4 = 1GE
- L5, L6 = 1GE
- L7, L8 = 1GE
- L9, L10 = 1GE
- L11, L12 = 10GE
- L13, L14 = 10GE

The 1GE links are more likely to be congested. In this example, the odd 1GE links have the following available bandwidth:

- L3 = 41%
- L5 = 56%
- L7 = 66%
- L9 = 71%

The even IGE links have the following available bandwidth:

- L4 = 37%
- L6 = 52%
- L8 = 61%
- L10 = 70%

Based on this information, the router would calculate the difference in available bandwidth between the odd links and the even links as follows:

- $L4 - L3 = 41\% - 37\% = 4\%$
- $L6 - L5 = 56\% - 52\% = 4\%$
- $L8 - L7 = 66\% - 61\% = 5\%$
- $L10 - L9 = 71\% - 70\% = 1\%$

The total additional bandwidth available over the odd links is 14% ($4\% + 4\% + 5\% + 1\%$). Since 14% is greater than 10% (the least-fill algorithm minimum threshold), the LSP is moved to the new path over the odd links from the original path using the even links.

8. Otherwise, the new path is rejected.

You can disable the following reoptimization criteria (a subset of the criteria listed previously):

- If the new path has the same IGP metric, it is not more hops away.
- The new path does not cause preemption. (This is to reduce the ripple effect of preemption causing more preemption.)
- The new path does not worsen congestion overall.
- If the new path has an equal IGP metric and lower hop count, it is accepted.

To disable them, either issue the **clear mpls lsp optimize-aggressive** command or include the **optimize-aggressive** statement:

optimize-aggressive;

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

Including the **optimize-aggressive** statement in the configuration causes the reoptimization procedure to be triggered more often. Paths are rerouted more frequently. It also limits the reoptimization algorithm to the IGP metric only.

**Related
Documentation**

- [Configuring Adaptive LSPs on page 249](#)
- [Configuring Automatic Bandwidth Allocation for LSPs on page 257](#)

- [Configuring MPLS Soft Preemption on page 238](#)
- [Configuring the Smart Optimize Timer for LSPs on page 255](#)
- [Configuring Hot Standby of Secondary Paths for LSPs on page 267](#)

Configuring the Smart Optimize Timer for LSPs

Because of network and router resource constraints, it is typically inadvisable to configure a short interval for the optimize timer. However, under certain circumstances, it might be desirable to reoptimize a path sooner than would normally be provided by the optimize timer.

For example, an LSP is traversing a preferred path that subsequently fails. The LSP is then switched to a less desirable path to reach the same destination. Even if the original path is quickly restored, it could take an excessively long time for the LSP to use it again, because it has to wait for the optimize timer to reoptimize the network paths. For such situations, you might want to configure the smart optimize timer.

When you enable the smart optimize timer, an LSP is switched back to its original path so long as the original path has been restored within 3 minutes of going down. Also, if the original path goes down again within 60 minutes, the smart optimize timer is disabled, and path optimization behaves as it normally does when the optimize timer alone is enabled. This prevents the router from using a flapping link.

The smart optimize timer is dependant on other MPLS features to function properly. For the scenario described here in which an LSP is switched to an alternate path in the event of a failure on the original path, it is assumed that you have configured one or more of the MPLS traffic protection features, including fast reroute, link protection, and standby secondary paths. These features help to ensure that traffic can reach its destination in the event of a failure.

At the least, you must configure a standby secondary path for the smart optimize timer feature to work properly. Fast reroute and link protection are more temporary solutions to a network outage. A secondary path ensures that there is a stable alternate path in the event the primary path fails. If you have not configured any sort of traffic protection for an LSP, the smart optimize timer by itself does not ensure that traffic can reach its destination. For more information about MPLS traffic protection, see [“MPLS and Traffic Protection” on page 45](#).

When a primary path fails and the smart optimize timer switches traffic to the secondary path, the router might continue to use the secondary path even after the primary path has been restored. If the ingress router completes a CSPF calculation, it might determine that the secondary path is the better path.

This might be undesirable if the primary path should be the active path and the secondary path should be used as a backup only. Also, if the secondary path is being used as the active path (even though the primary path has been reestablished) and the secondary path fails, the smart optimize timer feature will not automatically switch traffic back to the primary path. However, you can enable protection for the secondary path by

configuring node and link protection or an additional standby secondary path, in which case, the smart optimize timer can be effective.

Specify the time in seconds for the smart optimize timer using the **smart-optimize-timer** statement:

smart-optimize-timer *seconds*;

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

**Related
Documentation**

- [MPLS and Traffic Protection on page 45](#)
- [Optimizing Signaled LSPs on page 251](#)

Limiting the Number of Hops in LSPs

By default, each LSP can traverse a maximum of 255 hops, including the ingress and egress routers. To modify this value, include the **hop-limit** statement:

hop-limit *number*;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The number of hops can be from 2 through 255. (A path with two hops consists of the ingress and egress routers only.)

Configuring the Bandwidth Value for LSPs

Each LSP has a bandwidth value. This value is included in the sender's Tspec field in RSVP path setup messages. You can specify a bandwidth value in bits per second. If you configure more bandwidth for an LSP, it should be able to carry a greater volume of traffic. The default bandwidth is 0 bits per second.

A nonzero bandwidth requires that transit and egress routers reserve capacity along the outbound links for the path. The RSVP reservation scheme is used to reserve this capacity. Any failure in bandwidth reservation (such as failures at RSVP policy control or admission control) might cause the LSP setup to fail. If there is insufficient bandwidth on the interfaces for the transit or egress routers, the LSP is not established.

To specify a bandwidth value for a signaled LSP, include the **bandwidth** statement:

bandwidth *bps*;

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Automatic Bandwidth Allocation for LSPs

Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure an LSP with minimal bandwidth; this feature can dynamically adjust the LSP's bandwidth allocation based on current traffic patterns. The bandwidth adjustments do not interrupt traffic flow through the tunnel.

You set a sampling interval on an LSP configured with automatic bandwidth allocation. The average bandwidth is monitored during this interval. At the end of the interval, an attempt is made to signal a new path for the LSP with the bandwidth allocation set to the maximum average value for the preceding sampling interval. If the new path is successfully established and the original path is removed, the LSP is switched over to the new path. If a new path is not created, the LSP continues to use its current path until the end of the next sampling interval, when another attempt is made to establish a new path. Note that you can set minimum and maximum bandwidth values for the LSP.

During the automatic bandwidth allocation interval, the router might receive a steady increase in traffic (increasing bandwidth utilization) on an LSP, potentially causing congestion or packet loss. To prevent this, you can define a second trigger to prematurely expire the automatic bandwidth adjustment timer before the end of the current adjustment interval.

Configuring Automatic Bandwidth Allocation for LSPs

Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure an LSP with minimal bandwidth, and this feature can dynamically adjust the LSP's bandwidth allocation based on current traffic patterns. The bandwidth adjustments do not interrupt traffic flow through the tunnel.

At the end of the automatic bandwidth allocation time interval, the current maximum average bandwidth usage is compared with the allocated bandwidth for the LSP. If the LSP needs more bandwidth, an attempt is made to set up a new path where bandwidth is equal to the current maximum average usage. If the attempt is successful, the LSP's traffic is routed through the new path and the old path is removed. If the attempt fails, the LSP continues to use its current path.

If you have configured link and node protection for the LSP and traffic has been switched to the bypass LSP, the automatic bandwidth allocation feature continues to operate and take bandwidth samples from the bypass LSP. For the first bandwidth adjustment cycle, the maximum average bandwidth usage taken from the original link and node-protected LSP is used to resignal the bypass LSP if more bandwidth is needed. (Link and node protection are not supported on QFX Series switches.)

If you have configured fast-reroute for the LSP, you might not be able to use this feature to adjust the bandwidth. Because the LSPs use a fixed filter (FF) reservation style, when a new path is signaled, the bandwidth might be double-counted. Double-counting can

prevent a fast-reroute LSP from ever adjusting its bandwidth when automatic bandwidth allocation is enabled. (Fast reroute is not supported on QFX Series switches.)

To configure automatic bandwidth allocation, complete the steps in the following sections:

- [Configuring Automatic Bandwidth Allocation on LSPs on page 258](#)
- [Requesting Automatic Bandwidth Allocation Adjustment on page 263](#)

Configuring Automatic Bandwidth Allocation on LSPs

To enable automatic bandwidth allocation on an LSP, include the **auto-bandwidth** statement:

```
auto-bandwidth {  
  adjust-interval seconds;  
  adjust-threshold percent;  
  adjust-threshold-overflow-limit number;  
  adjust-threshold-underflow-limit number;  
  maximum-bandwidth bps;  
  minimum-bandwidth bps;  
  minimum-bandwidth-adjust-interval  
  minimum-bandwidth-adjust-threshold-change  
  minimum-bandwidth-adjust-threshold-value  
  monitor-bandwidth;  
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls *label-switched-path* *lsp-name*]**
- **[edit logical-systems *logical-system-name* protocols mpls *label-switched-path* *lsp-name*]**

If an LSP has an automatic bandwidth configuration, you can disable automatic bandwidth adjustments on a particular path (either primary or secondary) by configuring a static bandwidth value and by disabling the CSPF computation (using the **no-cspf** statement).

For example:

```
user@host> show protocols mpls  
label-switched-path primary-path {  
  to 192.168.0.1;  
  ldp-tunneling;  
  optimize-timer 3571;  
  least-fill;  
  link-protection;  
  adaptive;  
  auto-bandwidth {  
    adjust-interval 7177;  
    adjust-threshold 5;  
    minimum-bandwidth 1m;  
    maximum-bandwidth 2500000000;  
    adjust-threshold-overflow-limit 2;  
    resignal-minimum-bandwidth;  
  }  
  primary primary-path;
```



```

secondary secondary-path {
  bandwidth 0;
  no-cspf;
  priority 0 0;
}

```

The statements configured at the `[edit protocols mpls label-switched-path label-switched-path-name auto-bandwidth]` hierarchy level are optional and explained in the following sections:

- [Configuring the Automatic Bandwidth Allocation Interval on page 259](#)
- [Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 259](#)
- [Configuring the Automatic Bandwidth Adjustment Threshold on page 260](#)
- [Configuring a Limit on Bandwidth Overflow and Underflow Samples on page 261](#)
- [Configuring Passive Bandwidth Utilization Monitoring on page 263](#)

Configuring the Automatic Bandwidth Allocation Interval

At the end of the automatic bandwidth allocation interval, the automatic bandwidth computation and new path setup process is triggered.



NOTE: To prevent unnecessary resignaling of LSPs, it is best to configure an LSP adjustment interval that is at least three times longer than the MPLS automatic bandwidth statistics interval. For example, if you configure a value of 30 seconds for the MPLS automatic bandwidth statistics interval (interval statement at the `[edit protocols mpls statistics]` hierarchy level), you should configure a value of at least 90 seconds for the LSP adjustment interval (adjust-interval statement at the `[edit protocols mpls label-switched-path label-switched-path-name auto-bandwidth]` hierarchy level). See also “[Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs](#)” on page 264.

To specify the bandwidth reallocation interval in seconds for a specific LSP, include the `adjust-interval` statement:

```
adjust-interval seconds;
```

You can include this statement at the following hierarchy levels:

- `[edit protocols mpls label-switched-path lsp-name auto-bandwidth]`
- `[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name auto-bandwidth]`

Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth

You can maintain the LSP's bandwidth between minimum and maximum bounds by specifying values for the `minimum-bandwidth` and `maximum-bandwidth` statements.

To specify the minimum amount of bandwidth allocated for a specific LSP, include the **minimum-bandwidth** statement:

```
minimum-bandwidth bps;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name* **auto-bandwidth**]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name* **auto-bandwidth**]

To specify the maximum amount of bandwidth allocated for a specific LSP, include the **maximum-bandwidth** statement:

```
maximum-bandwidth bps;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name* **auto-bandwidth**]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name* **auto-bandwidth**]

Configuring the Automatic Bandwidth Adjustment Threshold

Use the **adjust-threshold** statement to specify the sensitivity of the automatic bandwidth adjustment of an LSP to changes in bandwidth utilization. You can set the threshold for when to trigger automatic bandwidth adjustments. When configured, bandwidth demand for the current interval is determined and compared to the LSP's current bandwidth allocation. If the percentage difference in bandwidth is greater than or equal to the specified **adjust-threshold** percentage, the LSP's bandwidth is adjusted to the current bandwidth demand.

For example, assume that the current bandwidth allocation is 100 megabits per second (Mbps) and that the percentage configured for the **adjust-threshold** statement is 15 percent. If the bandwidth demand increases to 110 Mbps, the bandwidth allocation is not adjusted. However, if the bandwidth demand increases to 120 Mbps (20 percent over the current allocation) or decreases to 80 Mbps (20 percent under the current allocation), the bandwidth allocation is increased to 120 Mbps or decreased to 80 Mbps, respectively.

To configure the threshold for automatic bandwidth adjustment, include the **adjust-threshold** statement:

```
adjust-threshold percent;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name* **auto-bandwidth**]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name* **auto-bandwidth**]

Configuring a Limit on Bandwidth Overflow and Underflow Samples

The automatic bandwidth adjustment timer is a periodic timer which is triggered every adjust interval to determine whether any bandwidth adjustments are required on the LSP's active path. This interval is typically configured as a long period of time, usually hours. If, at the end of adjust interval, the change in bandwidth is above a certain adjust threshold, the LSP is resigned with the new bandwidth.

During the automatic bandwidth adjustment interval, the router might receive a steady increase in traffic (increasing bandwidth utilization) on an LSP, potentially causing congestion or packet loss. To prevent this, you can define a second trigger to prematurely expire the automatic bandwidth adjustment timer before the end of the current adjustment interval.

Every statistics interval, the router samples the average bandwidth utilization of an LSP and if this has exceeded the current maximum average bandwidth utilization, the maximum average bandwidth utilization is updated.

During each sample period, the following conditions are also checked:

- Is the current average bandwidth utilization above the active bandwidth of the path?
- Has the difference between the average bandwidth utilization and the active bandwidth exceeded the adjust threshold (bandwidth utilization has changed significantly)?

If these conditions are true, it is considered to be one bandwidth overflow sample. Using the **adjust-threshold-overflow-limit** statement, you can define a limit on the number of bandwidth overflow samples such that when the limit is reached, the current automatic bandwidth adjustment timer is expired and a bandwidth adjustment is triggered. Once this adjustment is complete, the normal automatic bandwidth adjustment timer is reset to expire after the periodic adjustment interval.

To specify a limit on the number of bandwidth overflow samples before triggering an automatic bandwidth allocation adjustment, configure the **adjust-threshold-overflow-limit** statement:

```
adjust-threshold-overflow-limit number;
```

Similarly, if the current average bandwidth utilization is below the active bandwidth of the path by the configured adjusted threshold (meaning that bandwidth utilization has gone down significantly), the sample is considered to be an underflow sample. The adjusted (new signaling) bandwidth after an adjustment due to underflow is the maximum average bandwidth among the underflow samples. You can specify a limit on the number of bandwidth underflow samples before triggering an automatic bandwidth allocation adjustment by configuring the **adjust threshold-underflow-limit** statement:

```
adjust-threshold-underflow-limit number;
```

These statements can be configured at the following hierarchy levels:

- **[edit protocols mpls label-switched-path lsp-name auto-bandwidth]**

- [edit logical-systems *logical-system-name* protocols mpls [label-switched-path](#) *lsp-name* [auto-bandwidth](#)]

You must configure the **adjust-threshold** and **minimum-bandwidth** statements whenever you configure the **adjust-threshold-underflow-limit** statement. You must configure the **adjust-threshold** and **maximum-bandwidth** statements whenever you configure the **adjust-threshold-overflow-limit** statement

- You must configure a nonzero value for the **adjust-threshold** statement if you configure the **adjust-threshold-overflow-limit** or **adjust-threshold-underflow-limit** statement.
- Any bandwidth increase or decrease below the value configured for the **adjust-threshold** statement does not constitute an overflow or underflow condition.
- To prevent unlimited increases in LSP bandwidth (to limit overflow beyond a certain bandwidth), you must also configure the **maximum-bandwidth** statement when you configure the **adjust-threshold-overflow-limit** statement.

The following describes the other aspects of the **adjust-threshold-overflow-limit** statement:

- It only applies to bandwidth overflows. If the bandwidth is decreasing, the normal automatic bandwidth adjustment interval is used.
- It does not affect manually triggered automatic bandwidth adjustment.
- It applies to single-class DiffServ-TE LSPs.
- Because the **adjust-threshold-overflow-limit** statement can trigger a bandwidth adjustment, it cannot be enabled at the same time as the **monitor-bandwidth** statement (for information about that statement, see [“Configuring Passive Bandwidth Utilization Monitoring” on page 263](#)).
- You cannot configure automatic bandwidth adjustments to occur more often than every 300 seconds. The **adjust-threshold-overflow-limit** statement is subject to the same minimum value with regard to the minimum frequency of adjustment allowed. Overflow condition based adjustments can occur no sooner than 300 seconds from the start of the overflow condition. Therefore it is required that:

sample interval x adjust-threshold-overflow-limit >= 300s

These values are checked during the commit operation. An error is returned if the value is less than 300 seconds.

- If you change the value of the **adjust-threshold-overflow-limit** statement on a working router, you can expect the following behavior:
 - If you increase the current value of the **adjust-threshold-overflow-limit** statement, the old value is replaced with the new one.
 - If you decrease the current value of the **adjust-threshold-overflow-limit** statement and the current bandwidth overflow count is less than the new value, the old value is replaced with the new one.

- If you decrease the current value of the **adjust-threshold-overflow-limit** statement and the current bandwidth overflow count is greater than the new value, the adjustment timer is immediately expired and a bandwidth adjustment is initiated.

Configuring Passive Bandwidth Utilization Monitoring

Use the **monitor-bandwidth** statement to switch to a passive bandwidth utilization monitoring mode. In this mode, no automatic bandwidth adjustments are made, but the maximum average bandwidth utilization is continuously monitored and recorded.

To configure passive bandwidth utilization monitoring, include the **monitor-bandwidth** statement:

```
monitor-bandwidth;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name* **auto-bandwidth**]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name* **auto-bandwidth**]

If you have configured an LSP with primary and secondary paths, the automatic bandwidth allocation statistics are carried over to the secondary path if the primary path fails. For example, consider a primary path whose adjustment interval is half complete and whose maximum average bandwidth usage is currently calculated as 50 Mbps. If the primary path suddenly fails, the time remaining for the next adjustment and the maximum average bandwidth usage are carried over to the secondary path.

Requesting Automatic Bandwidth Allocation Adjustment

For MPLS LSP automatic bandwidth allocation adjustment, the minimum value for the adjustment interval is 5 minutes (300 seconds). You might find it necessary to trigger a bandwidth allocation adjustment manually, for example in the following circumstances:

- When you are testing automatic bandwidth allocation in a network lab.
- When the LSP is configured for automatic bandwidth allocation in monitor mode (the **monitor-bandwidth** statement is included in the configuration as described in [“Configuring Passive Bandwidth Utilization Monitoring” on page 263](#)), and want to initiate an immediate bandwidth adjustment.

To use the **request mpls lsp adjust-autobandwidth** command, the following must be true:

- Automatic bandwidth allocation must be enabled on the LSP.
- The criteria required to trigger a bandwidth adjustment have been met (the difference between the adjust bandwidth and the current LSP path bandwidth is greater than the threshold limit).

A manually triggered bandwidth adjustment operates only on the active LSP path. Also, if you have enabled periodic automatic bandwidth adjustment, the periodic automatic bandwidth adjustment parameters (the adjustment interval and the maximum average bandwidth) are not reset after a manual adjustment.

For example, suppose the periodic adjust interval is 10 hours and there are currently 5 hours remaining before an automatic bandwidth adjustment is triggered. If you initiate a manual adjustment with the **request mpls lsp adjust-autobandwidth** command, the adjust timer is not reset and still has 5 hours remaining.

To manually trigger a bandwidth allocation adjustment, you need to use the **request mpls lsp adjust-autobandwidth** command. You can trigger the command for all affected LSPs on the router, or you can specify a particular LSP:

```
user@host> request mpls lsp adjust-autobandwidth
```

Once you execute this command, the automatic bandwidth adjustment validation process is triggered. If all the criteria for adjustment are met, the LSP's active path bandwidth is adjusted to the adjusted bandwidth value determined during the validation process.

**Related
Documentation**

- [Configuring MPLS to Gather Statistics on page 342](#)
- [Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs on page 264](#)
- [request mpls lsp adjust-autobandwidth on page 1168](#)
- [show mpls lsp on page 1209](#)

Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs

Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure the device to collect statistics related to automatic bandwidth allocation by completing the following steps:

1. To collect statistics related to automatic bandwidth allocation, configure the **auto-bandwidth** option for the **statistics** statement at the **[edit protocols mpls]** hierarchy level. These settings apply to all LSPs configured on the router on which you have also configured the **auto-bandwidth** statement at the **[edit protocols mpls label-switched-path label-switched-path-name]** hierarchy level.

```
statistics {  
  auto-bandwidth;  
  file filename <files number> <size size> <world-readable | no-world-readable>;  
  interval seconds;  
  no-transit-statistics;  
  transit-statistics-polling;  
}
```

2. Specify the **filename** for the files used to store the MPLS trace operation output using the **file** option. All files are placed in the directory **/var/log**. We recommend that you place MPLS tracing output in the file **mpls-log**.
3. Specify the maximum number of trace files using the **files number** option. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.
4. Specify the interval for calculating the average bandwidth usage by configuring a time in seconds using the **interval** option. You can also set the adjustment interval on a

specific LSP by configuring the **interval** option at the **[edit protocols mpls label-switch-path label-switched-path-name statistics]** hierarchy level.



NOTE: To prevent unnecessary resignaling of LSPs, it is best to configure an LSP adjustment interval that is at least three times longer than the MPLS automatic bandwidth statistics interval. For example, if you configure a value of 30 seconds for the MPLS automatic bandwidth statistics interval (interval statement at the **[edit protocols mpls statistics]** hierarchy level), you should configure a value of at least 90 seconds for the LSP adjustment interval (adjust-interval statement at the **[edit protocols mpls label-switched-path label-switched-path-name auto-bandwidth]** hierarchy level).

5. To trace automatic bandwidth allocation, include the **autobw-state** flag for the MPLS **traceoptions** statement at the **[edit protocols mpls]** hierarchy level.

The following configuration enables the MPLS traceoptions for automatic bandwidth allocation. The trace records are stored in a file called **auto-band-trace** (the filename is user configurable):

```
[edit protocols mpls]
traceoptions {
    file auto-band-trace size 10k files 10 world-readable;
    flag autobw-state;
}
```

6. Using the **show log** command, you can display the automatic bandwidth allocation statistics file generated when you configure the **auto-bandwidth** statement. The following shows sample log file output taken from an MPLS statistics file named **auto-band-stats** on a router configured with an LSP named **E-D**. The log file shows that LSP **E-D** is operating over its reserved bandwidth limit initially. Before **Oct 30 17:14:57**, the router triggered an automatic bandwidth adjustment (you might see two sessions for an LSP undergoing an automatic bandwidth adjustment). By **Oct 30 17:16:57**, the LSP has been reestablished at a higher bandwidth and is now shown using less than 100 percent of its **Reserved Bw** (reserved bandwidth).

```
user@host> show log auto-band-stats
E-D          (LSP ID 5, Tunnel ID 6741)          209 pkt          17094 Byte          1 pps          90 Bps Util
 240.01% Reserved Bw          37 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10Oct 30 17:13:57 Total 1 sessions: 1
success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)          241 pkt          19737 Byte          1 pps          88 Bps Util
 234.67% Reserved Bw          37 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10Oct 30 17:14:27 Total 1 sessions: 1
success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)          276 pkt          22607 Byte          1 pps          95 Bps Util
 253.34% Reserved Bw          37 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10Oct 30 17:14:57 Total 1 sessions: 1
success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)           0 pkt              0 Byte            0 pps           0 Bps Util
  0.00% Reserved Bw          37 Bps
E-D          (LSP ID 6, Tunnel ID 6741)           0 pkt              0 Byte            0 pps           0 Bps Util
  0.00% Reserved Bw          101 Bps
```

```

decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh 0x952c308, type 4, flags 0x0,
n_gw 1, nhid 0 to refcount 1Oct 30 17:15:27 Total 2 sessions: 2 success, 0 fail, 0 ignored
E-D      (LSP ID 5, Tunnel ID 6741)      0 pkt      0 Byte      0 pps      0 Bps Util
  0.00% Reserved Bw      37 Bps
E-D      (LSP ID 6, Tunnel ID 6741)      33 pkt      2695 Byte      1 pps      89 Bps Util
  87.69% Reserved Bw      101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh 0x952c308, type 4, flags 0x0,
n_gw 1, nhid 0 to refcount 1Oct 30 17:15:57 Total 2 sessions: 2 success, 0 fail, 0 ignored
E-D      (LSP ID 5, Tunnel ID 6741)      0 pkt      0 Byte      0 pps      0 Bps Util
  0.00% Reserved Bw      37 Bps
E-D      (LSP ID 6, Tunnel ID 6741)      65 pkt      5338 Byte      1 pps      88 Bps Util
  86.70% Reserved Bw      101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh 0x952c308, type 4, flags 0x0,
n_gw 1, nhid 0 to refcount 1Oct 30 17:16:27 Total 2 sessions: 2 success, 0 fail, 0 ignored
E-D      (LSP ID 6, Tunnel ID 6741)      97 pkt      7981 Byte      1 pps      88 Bps Util
  86.70% Reserved Bw      101 Bps
decr nh 0x952c308, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1Oct 30 17:16:57 Total 1 sessions: 1
success, 0 fail, 0 ignored

```

7. Issue the `show mpls lsp autobandwidth` command to display current information about automatic bandwidth allocation. The following shows sample output from the `show mpls lsp autobandwidth` command taken at about the same time as the log file shown previously:

```

user@host> show mpls lsp autobandwidth
Lspname      Last      Requested      Reserved      Highwater      AdjustTime LastAdjust
BW           BW           BW           mark           Left (sec)
E-D          300bps      812.005bps    812bps        1.56801kbps    294 sec      Wed Oct 30 17:15:26 2013

```

8. Issue the `file show` command to display the MPLS trace file. You need to specify the file location and file name (the file is located in `/var/log/`). The following shows sample trace file output is taken from an MPLS trace file named `auto-band-trace.0.gz` on a router configured with an LSP named `E-D`. The trace file shows that LSP `E-D` is operating over its reserved bandwidth limit initially. At **Oct 30 17:15:26**, the router triggers an automatic bandwidth adjustment (you might see two sessions for an LSP undergoing an automatic bandwidth adjustment). By **Oct 30 17:15:57**, the LSP has been reestablished at a higher bandwidth and is now shown using less than 100 percent of its **Reserved Bw** (reserved bandwidth).

```

user@host> file show /var/log/auto-band-trace.0.gz
Oct 30 17:13:57 trace_on: Tracing to "/var/log/E/auto-band-trace" started
Oct 30 17:13:57.466825 LSP E-D (id 5) new bytes arrived      2714 in 29
sec
Oct 30 17:14:27.466713 E-D      (LSP ID 5, Tunnel ID 6741)      241
pkt      19737 Byte      1 pps      88 Bps Util 234.67% Reserved Bw
      37 Bps
Oct 30 17:14:27.466962 LSP E-D (id 5, old id 5); sampled bytes      19737 >
bytes recorded      17094
Oct 30 17:14:27.467035 LSP E-D (id 5) new bytes arrived      2643 in 29
sec
Oct 30 17:14:57.466599 E-D      (LSP ID 5, Tunnel ID 6741)      276
pkt      22607 Byte      1 pps      95 Bps Util 253.34% Reserved Bw
      37 Bps
Oct 30 17:14:57.466758 LSP E-D (id 5, old id 5); sampled bytes      22607 >
bytes recorded      19737
Oct 30 17:14:57.466825 LSP E-D (id 5) new bytes arrived      2870 in 29
sec
Oct 30 17:15:26.265816 Adjust Autobw: LSP E-D (id 5) curr adj bw 300bps updated
with 812.005bps

```



```

Oct 30 17:15:26.266064 mpls LSP E-D Autobw change 512.005bps >= threshold 75bps
Oct 30 17:15:26.363372 Autobw Success: LSP E-D () (old id 5 new id 6) update
  prev active bw 300 bps with 812 bps
Oct 30 17:15:26.363686 RPD_MPLS_PATH_BANDWIDTH_CHANGE: MPLS path (lsp E-D)
bandwidth changed, path bandwidth 812 bps
Oct 30 17:15:27.364751 RPD_MPLS_LSP_BANDWIDTH_CHANGE: MPLS LSP E-D bandwidth
changed, lsp bandwidth 812 bps
Oct 30 17:15:27.466849 E-D (LSP ID 5, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
37 Bps
Oct 30 17:15:27.467050 E-D (LSP ID 6, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
101 Bps
Oct 30 17:15:57.466858 E-D (LSP ID 5, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
37 Bps
Oct 30 17:15:57.467106 E-D (LSP ID 6, Tunnel ID 6741) 33
pkt 2695 Byte 1 pps 89 Bps Util 87.69% Reserved Bw
101 Bps
Oct 30 17:15:57.467201 LSP E-D (id 6, old id 5); LSP up after autobw adjustment
and active for 30 sec
Oct 30 17:15:57.467398 LSP E-D (id 6) psb bytes 2695 < bytes recorded
22607 total bytes 2695 in 30 sec
Oct 30 17:15:57.467461 First sample of the adjust interval after automatic bw
adjustment
Oct 30 17:15:57.467594 Update curr max avg bw 0bps of LSP E-D with new bw
716.225bps
Oct 30 17:16:27.466830 E-D (LSP ID 5, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
37 Bps
Oct 30 17:16:27.467079 E-D (LSP ID 6, Tunnel ID 6741) 65
pkt 5338 Byte 1 pps 88 Bps Util 86.70% Reserved Bw
101 Bps
Oct 30 17:16:27.467171 LSP E-D (id 6, old id 6); sampled bytes 5338 >
bytes recorded 2695
Oct 30 17:16:27.467237 LSP E-D (id 6) new bytes arrived 2643 in 29
sec
Oct 30 17:16:57.466712 E-D (LSP ID 6, Tunnel ID 6741) 97
pkt 7981 Byte 1 pps 88 Bps Util 86.70% Reserved Bw
101 Bps
Oct 30 17:16:57.466870 LSP E-D (id 6, old id 6); sampled bytes 7981 >
bytes recorded 5338

```

- Related Documentation**
- [Configuring Automatic Bandwidth Allocation for LSPs on page 257](#)
 - [show mpls lsp autobandwidth on page 1227](#)

Configuring Hot Standby of Secondary Paths for LSPs

By default, secondary paths are set up only as needed. To have the system maintain a secondary path in a hot-standby state indefinitely, include the **standby** statement:

standby;

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls label-switched-path *lsp-name* secondary]**

- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name* secondary]

The hot-standby state is meaningful only on secondary paths. Maintaining a path in a hot-standby state enables swift cutover to the secondary path when downstream routers on the current active path indicate connectivity problems. Although it is possible to configure the **standby** statement at the [edit protocols mpls label-switched-path *lsp-name* primary *path-name*] hierarchy level, it has no effect on router behavior.

If you configure the **standby** statement at the following hierarchy levels, the hot-standby state is activated on all secondary paths configured beneath that hierarchy level:

- [edit protocols mpls]
- [edit protocols mpls label-switched-path *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]

The hot-standby state has two advantages:

- It eliminates the call-setup delay during network topology changes. Call setup can suffer from significant delays when network failures trigger large numbers of LSP reroutes at the same time.
- A cutover to the secondary path can be made before RSVP learns that an LSP is down. There can be significant delays between the time the first failure is detected by protocol machinery (which can be an interface down, a neighbor becoming unreachable, a route becoming unreachable, or a transient routing loop being detected) and the time an LSP actually fails (which requires a timeout of soft state information between adjacent RSVP routers). When topology failures occur, hot-standby secondary paths can usually achieve the smallest cutover delays with minimal disruptions to user traffic.

When the primary path is considered to be stable again, traffic is automatically switched from the standby secondary path back to the primary path. The switch is performed no faster than twice the retry-timer interval and only if the primary path exhibits stability throughout the entire switch interval.

The drawback of the hot-standby state is that more state information must be maintained by all the routers along the path, which requires overhead from each of the routers.

Damping Advertisement of LSP State Changes

When an LSP changes from being up to being down, or from down to up, this transition takes effect immediately in the router software and hardware. However, when advertising LSPs into IS-IS and OSPF, you may want to damp LSP transitions, thereby not advertising the transition until a certain period of time has transpired (known as the hold time). In this case, if the LSP goes from up to down, the LSP is not advertised as being down until it has remained down for the hold-time period. Transitions from down to up are advertised into IS-IS and OSPF immediately. Note that LSP damping affects only the IS-IS and OSPF advertisements of the LSP; other routing software and hardware react immediately to LSP transitions.

To damp LSP transitions, include the **advertisement-hold-time** statement:

advertisement-hold-time *seconds*;

seconds can be a value from 0 through 65,535 seconds. The default is 5 seconds.

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

CHAPTER 5

Configuring Static and Explicit-Path LSPs

- [Configuring Static LSPs on page 271](#)
- [Configuring Explicit-Path LSPs on page 278](#)

Configuring Static LSPs

To configure static LSPs, configure the ingress router and each router along the path up to and including the egress router.

To configure static MPLS, perform the following tasks:

- [Configuring the Ingress Router for Static LSPs on page 271](#)
- [Configuring the Intermediate \(Transit\) and Egress Routers for Static LSPs on page 274](#)
- [Configuring a Bypass LSP for the Static LSP on page 276](#)
- [Configuring the Protection Revert Timer for Static LSPs on page 277](#)
- [Configuring Static Unicast Routes for Point-to-Multipoint LSPs on page 277](#)

Configuring the Ingress Router for Static LSPs

The ingress router checks the IP address in the incoming packet's destination address field and, if it finds a match in the routing table, applies the label associated with that address to the packets. The label has forwarding information associated with it, including the address of the next-hop router, and the route preference and CoS values.

To configure static LSPs on the ingress router, include the **ingress** statement:

```
ingress {  
  bandwidth bps;  
  class-of-service cos-value;  
  description string;  
  install {  
    destination-prefix <active>;  
  }  
  link-protection bypass-name name;  
  metric metric;  
  next-hop (address | interface-name | address/interface-name);  
  no-install-to-address;  
  node-protection bypass-name name next-next-label label;  
  policing {  
    filter filter-name;
```

```
        no-auto-policing;
    }
    preference preference;
    push out-label;
    to address;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols mpls static-label-switched-path *static-lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *static-lsp-name*]

When you configure a static LSP on the ingress router, the **next-hop**, **push**, and **to** statements are required; the other statements are optional.

The configuration for a static LSP on the ingress router requires you to configure the following parts:

- Criteria for analyzing an incoming packet:
 - The **install** statement creates an LSP that handles IPv4 packets. All static MPLS routes created using the **install** statement are installed in inet.3 routing table, and the creating protocol is identified as static. This process is no different from creating static IPv4 routes at the [edit routing-options static] hierarchy level.
 - In the **to** statement, you configure the IP destination address to check when incoming packets are analyzed. If the address matches, the specified outgoing label (**push out-label**) is assigned to the packet, and the packet enters an LSP. Manually assigned outgoing labels can have values from 0 through 1,048,575. Each prefix that you specify is installed as a static route in the routing table.
- The **next-hop** statement, which supplies the IP address of the next hop to the destination. You can specify this as the IP address of the next hop, the interface name (for point-to-point interfaces only), or as **address/interface-name** to specify an IP address on an operational interface. When the next hop is on a directly attached interface, the route is installed in the routing table. You cannot configure a LAN or nonbroadcast multiaccess (NBMA) interface as a next-hop interface.
- Properties to apply to the LSP (all are optional):
 - Bandwidth reserved for this LSP (**bandwidth bps**)
 - Link protection and node protection to apply to the LSP (**bypass bypass-name**, **link-protection bypass-name name**, **node-protection bypass-name next-next-label label**)
 - Metric value to apply to the LSP (**metric**)
 - Class-of-service value to apply to the LSP (**class-of-service**)
 - Preference value to apply to the LSP (**preference**)
 - Traffic policing to apply to the LSP (**policing**)

- Text description to apply to the LSP ([description](#))
- Install or no-install policy ([install](#) or [no-install-to-address](#))


To determine whether a static ingress route is installed, use the command **show route table inet.3 protocol static**. Sample output follows. The **push** keyword denotes that a label is to be added in front of an IP packet.

```
10.0.0.0      *[Static/5] 00:01:48
> to 11.1.1.1 via so-0/0/0, push 1000123
```

Example: Configuring the Ingress Router

Configure the ingress router for a static LSP that consists of three routers (see [Figure 31 on page 273](#)).

Figure 31: Static MPLS Configuration



For packets addressed to 10.0.0.0, assign label 1000123 and transmit them to the next-hop router at 11.1.1.1:

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
protocols {
  mpls {
    static-label-switched-path path1 {
      ingress {
        next-hop 11.1.1.1;
        to 10.0.0.0;
        push 1000123;
      }
    }
  }
  interface so-0/0/0.0;
}
}
routing-options {
  static {
    route 10.0.0.0/8 {
```

```

        static-lsp-next-hop path];
    }
}

```

To determine whether the static ingress route is installed, use the following command:

```
user@host> show route table inet.0 protocol static
```

Sample output follows. The **push 1000123** keyword identifies the route.

```

10.0.0.0/8          *[Static/5] 00:01:48
> to 11.1.1.1 via so-0/0/0.0, push 1000123

```

Configuring the Intermediate (Transit) and Egress Routers for Static LSPs

Intermediate (transit) and egress routers perform similar functions—they modify the label that has been applied to a packet. An intermediate router can change the label. An egress router removes the label (if the packet still contains a label) and continues forwarding the packet to its destination.

To configure static LSPs on intermediate and egress routers, include the **transit** statement:

```

static-label-switched-path lsp-name {
    transit incoming-label {
        bandwidth bps;
        description string;
        link-protection bypass-name name;
        next-hop (address | interface-name | address/interface-name);
        node-protection bypass-name name next-next-label label;
        pop;
        swap out-label;
    }
}

```

You can include these statements at the following hierarchy levels:

- [edit protocols mpls static-label-switched-path *static-lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *static-lsp-name*]

For the **transit** statement configuration, the **next-hop** and **pop** | **swap** statements are required. The remaining statements are optional.

Each statement within the **transit** statement consists of the following parts:

- Packet label (specified in the **transit** statement)
- The **next-hop** statement, which supplies the IP address of the next hop to the destination. The address is specified as the IP address of the next hop, or the interface name (for point-to-point interfaces only), or **address** and **interface-name** to specify an IP address on an operational interface. When the specified next hop is on a directly attached interface, this route is installed in the routing table. You cannot configure a LAN or NBMA interface as a next-hop interface.

- Operation to perform on the labeled packet:
 - For egress routers, you generally just remove the packet's label altogether (**pop**) and continue forwarding the packet to the next hop. However, if the previous router removed the label, the egress router examines the packet's IP header and forwards the packet toward its IP destination.
 - For intermediate (transit) routers only, exchange the label for another label (**swap out-label**). Manually assigned incoming labels can have values from 1,000,000 through 1,048,575. Manually assigned outgoing labels can have values from 0 through 1,048,575.
- Label properties to apply to the packet (all are optional):
 - Bandwidth reserved for this route (**bandwidth bps**).
 - Link-protection and node-protection to apply to the LSP (**bypass bypass-name, link-protection bypass-name name, node-protection bypass-name next-next-label label**).
 - Text description to apply to the LSP (specified in the **description** statement).

The static routes are installed in the default MPLS routing table, `mpls.0`, and the creating protocol is identified as static. To verify that a static route is properly installed, use the command **show route table mpls.0 protocol static**. Sample output follows:

```
1000123      *[Static/5] 00:00:38
> to 12.2.2.2 via so-5/0/0.0, swap 1000456
```

You can configure a revert timer for a static LSP transiting an intermediate router. After traffic has been switched to a bypass static LSP, it is typically switched back to the primary static LSP when it comes back up. There is a configurable delay in the time (called the revert timer) between when the primary static LSP comes up and when traffic is reverted back to it from the bypass static LSP. This delay is needed because when the primary LSP comes back up, it is not certain whether all of the interfaces on the downstream node of the primary path have come up yet. You can display the revert timer value for an interface using the **show mpls interface detail** command. For more information, see [“Configuring the Revert Timer for LSPs” on page 215](#).

Example: Configuring an Intermediate Router

For packets labeled **1000123** arriving on interface **so-0/0/0**, assign the label **1000456**, and transmit them to the next-hop router at **12.2.2.2**:

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
protocols {
  mpls {
    static-label-switched-path path1 {
```

```

        transit 1000123 {
            next-hop 12.2.2.2;
            swap 1000456;
        }
    }
    interface so-0/0/0.0;
}

```

To determine whether the static intermediate route is installed, use the following command:

```
user@host> show route table mpls.0 protocol static
```

Sample output follows. The **swap 1000456** keyword identifies the route.

```

1000123          *[Static/5] 00:01:48
> to 12.2.2.2 via so-0/0/0, swap 1000456

```

Example: Configuring an Egress Router

For packets labeled **1000456** arriving on interface **so-0/0/0**, remove the label and transmit the packets to the next-hop router at **13.3.3.3**:

```

[edit]
interfaces {
    so-0/0/0 {
        unit 0 {
            family mpls;
        }
    }
}
protocols {
    mpls {
        static-label-switched-path path1 {
            transit 1000456 {
                next-hop 13.3.3.3;
                pop;
            }
        }
        interface so-0/0/0.0;
    }
}

```

To determine whether the static egress route is installed, use the following command:

```
user@host> show route table mpls.0 protocol static
```

Sample output follows. The **pop** keyword identifies the egress route.

```

1000456          *[Static/5] 00:01:48
> to 13.3.3.3 via so-0/0/0, pop

```

Configuring a Bypass LSP for the Static LSP

To enable a bypass LSP for the static LSP, configure the **bypass** statement:

```

bypass bypass-name {
    bandwidth bps;
}

```

```

description string;
next-hop (address | interface-name | address/interface-name);
push out-label;
to address;
}

```

Configuring the Protection Revert Timer for Static LSPs

For static LSPs configured with a bypass static LSP, it is possible to configure the protection revert timer. If a static LSP goes down and traffic is switched to the bypass LSP, the protection revert timer specifies the amount of time (in seconds) that the LSP must wait before it can revert back to the original static LSP.

The range of values you can configure for the protection revert timer is 0 through 65,535 seconds. The default value is 5 seconds.

If you configure a value of 0 seconds, the traffic on the LSP, once switched from the original static LSP to the bypass static LSP, remains on the bypass LSP permanently (until the network operator intervenes or until the bypass LSP goes down).

You can configure the protection revert timer for all LSPs on the router at the **[edit protocols mpls]** hierarchy level or for a specific LSP at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level.

To configure the protection revert timer for static LSPs include the **protection-revert-time** statement:

```
protection-revert-time seconds;
```

For a list of hierarchy levels at which you can include this statement, see the summary section for this statement.

Configuring Static Unicast Routes for Point-to-Multipoint LSPs

You can configure a static unicast IP route with a point-to-multipoint LSP as the next hop. For more information about point-to-multipoint LSPs, see [“Point-to-Multipoint LSPs Overview” on page 281](#), [“Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs” on page 283](#), and [“Configuring CCC Switching for Point-to-Multipoint LSPs” on page 665](#).

To configure a static unicast route for a point-to-multipoint LSP, complete the following steps:

1. On the ingress PE router, configure a static IP unicast route with the point-to-multipoint LSP name as the next hop by including the **p2mp-lsp-next-hop** statement:

```
p2mp-lsp-next-hop point-to-multipoint-lsp-next-hop;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options static route *route-name*]**
- **[edit logical-systems *logical-system-name* routing-options static route *route-name*]**

2. On the egress PE router, configure a static IP unicast route with the same destination address configured in Step 1 (the address configured at the **[edit routing-options static route]** hierarchy level) by including the **next-hop** statement:

```
next-hop address;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options static route route-name]**
- **[edit logical-systems logical-system-name routing-options static route route-name]**



NOTE: CCC and static routes cannot use the same point-to-multipoint LSP.

For more information on static routes, see the *Junos OS Routing Protocols Library for Routing Devices*.

The following **show route** command output displays a unicast static route pointing to a point-to-multipoint LSP on the ingress PE router where the LSP has two branch next hops:

```
user@host> show route 5.5.5.5 detail
inet.0: 29 destinations, 30 routes (28 active, 0 holddown, 1 hidden)
5.5.5.5/32 (1 entry, 1 announced)
  *Static Preference: 5
    Next hop type: Flood
    Next hop: via so-0/3/2.0 weight 1
    Label operation: Push 100000
    Next hop: via t1-0/1/1.0 weight 1
    Label operation: Push 100064
    State: <Active Int Ext>
    Local AS: 10458
    Age: 2:41:15
    Task: RT
    Announcement bits (2): 0-KRT 3-BGP.0.0.0.0+179
    AS path: I
```

Configuring Explicit-Path LSPs

If you disable constrained-path label-switched path (LSP) computation, as described in [“Disabling Constrained-Path LSP Computation” on page 239](#), you can configure LSPs manually or allow the LSPs to follow the IGP path.

When explicit-path LSPs are configured, the LSP is established along the path you specified. If the path is topologically not feasible, either because the network is partitioned or insufficient resources are available along some parts of the path, the LSP will fail. No alternative paths can be used. If the setup succeeds, the LSP stays on the defined path indefinitely.

To configure an explicit-path LSP, follow these steps:

1. Configure the path information in a named path, as described in [“Creating Named Paths” on page 60](#). To configure complete path information, specify every router hop between the ingress and egress routers, preferably using the **strict** attribute. To configure incomplete path information, specify only a subset of router hops, using the **loose** attribute in places where the path is incomplete.

For incomplete paths, the MPLS routers complete the path by querying the local routing table. This query is done on a hop-by-hop basis, and each router can figure out only enough information to reach the next explicit hop. It might be necessary to traverse a number of routers to reach the next (loose) explicit hop.

Configuring incomplete path information creates portions of the path that depend on the current routing table, and this portion of the path can reroute itself as the topology changes. Therefore, an explicit-path LSP that contains incomplete path information is not completely fixed. These types of LSPs have only a limited ability to repair themselves, and they tend to create loops or flaps depending on the contents of the local routing table.

2. To configure the LSP and point it to the named path, use either the **primary** or **secondary** statement, as described in [“Configuring Primary and Secondary LSPs” on page 214](#).
3. Disable constrained-path LSP computation by including the **no-cspf** statement either as part of the LSP or as part of a **primary** or **secondary** statement. For more information, see [“Disabling Constrained-Path LSP Computation” on page 239](#).
4. Configure any other LSP properties.

Using explicit-path LSPs has the following drawbacks:

- More configuration effort is required.
- Configured path information cannot take into account dynamic network bandwidth reservation, so the LSPs tend to fail when resources become depleted.
- When an explicit-path LSP fails, you might need to manually repair it.

Because of these limitations, we recommend that you use explicit-path LSPs only in controlled situations, such as to enforce an optimized LSP placement strategy resulting from computations with an offline simulation software package.

CHAPTER 6

Configuring Point-to-Multipoint LSPs

- [Point-to-Multipoint LSPs Overview on page 281](#)
- [Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs on page 283](#)
- [Example: Configuring a Collection of Paths to Create an RSVP-Signaled Point-to-Multipoint LSP on page 285](#)
- [Configuring Inter-Domain Point-to-Multipoint LSPs on page 303](#)
- [Configuring Link Protection for Point-to-Multipoint LSPs on page 304](#)
- [Configuring Graceful Restart for Point-to-Multipoint LSPs on page 305](#)
- [Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs on page 306](#)
- [Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs on page 307](#)
- [Enabling Point-to-Point LSPs to Monitor Egress PE Routers on page 307](#)
- [Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases on page 308](#)


Point-to-Multipoint LSPs Overview

A point-to-multipoint MPLS LSP is an LSP with a single source and multiple destinations. By taking advantage of the MPLS packet replication capability of the network, point-to-multipoint LSPs avoid unnecessary packet replication at the ingress router. Packet replication takes place only when packets are forwarded to two or more different destinations requiring different network paths.

This process is illustrated in [Figure 32 on page 282](#). Router PE1 is configured with a point-to-multipoint LSP to Routers PE2, PE3, and PE4. When Router PE1 sends a packet on the point-to-multipoint LSP to Routers P1 and P2, Router P1 replicates the packet and forwards it to Routers PE2 and PE3. Router P2 sends the packet to Router PE4.

This feature is described in detail in the Internet drafts [draft-raggarwa-mpls-p2mp-te-02.txt](#) (expired February 2004), *Establishing Point to Multipoint MPLS TE LSPs*, [draft-ietf-mpls-rsvp-te-p2mp-02.txt](#), *Extensions to Resource Reservation Protocol-Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label-Switched Paths (LSPs)*, and RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths* (only point-to-multipoint LSPs are supported).

Figure 32: Point-to-Multipoint LSPs



The following are some of the properties of point-to-multipoint LSPs:

- A point-to-multipoint LSP enables you to use MPLS for point-to-multipoint data distribution. This functionality is similar to that provided by IP multicast.
- You can add and remove branch LSPs from a main point-to-multipoint LSP without disrupting traffic. The unaffected parts of the point-to-multipoint LSP continue to function normally.
- You can configure a node to be both a transit and an egress router for different branch LSPs of the same point-to-multipoint LSP.
- You can enable link protection on a point-to-multipoint LSP. Link protection can provide a bypass LSP for each of the branch LSPs that make up the point-to-multipoint LSP. If any of the primary paths fail, traffic can be quickly switched to the bypass.
- You can configure branch LSPs either statically, dynamically, or as a combination of static and dynamic LSPs.
- You can enable graceful Routing Engine switchover (GRES) and graceful restart for point-to-multipoint LSPs at ingress and egress routers. The point-to-multipoint LSPs must be configured using either static routes or circuit cross-connect (CCC). GRES and graceful restart allow the traffic to be forwarded at the Packet Forwarding Engine based on the old state while the control plane recovers. Feature parity for GRES and graceful restart for MPLS point-to-multipoint LSPs on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.

- Related Documentation**
- *Junos OS High Availability Library for Routing Devices*
 - *Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on Logical Systems*
 - *Example: NG-VPLS Using Point-to-Multipoint LSPs*
 - *Example: Configuring Point-to-Multipoint LDPLSPs as the Data Plane for Intra-AS MBGP MVPNs*

Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs

A point-to-multipoint MPLS label-switched path (LSP) is an RSVP LSP with multiple destinations. By taking advantage of the MPLS packet replication capability of the network, point-to-multipoint LSPs avoid unnecessary packet replication at the ingress router. For more information about point-to-multipoint LSPs, see [“Point-to-Multipoint LSPs Overview” on page 281](#).

To configure a point-to-multipoint LSP, you need to configure the primary LSP from the ingress router and the branch LSPs that carry traffic to the egress routers, as described in the following sections:

- [Configuring the Primary Point-to-Multipoint LSP on page 283](#)
- [Configuring a Branch LSP for Point-to-Multipoint LSPs on page 283](#)

Configuring the Primary Point-to-Multipoint LSP

A point-to-multipoint LSP must have a configured primary point-to-multipoint LSP to carry traffic from the ingress router. The configuration of the primary point-to-multipoint LSP is similar to a signaled LSP. See [“Configuring the Ingress Router for MPLS-Signaled LSPs” on page 60](#) for more information. In addition to the conventional LSP configuration, you need to specify a path name for the primary point-to-multipoint LSP by including the **p2mp** statement:

```
p2mp p2mp-lsp-name;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls label-switched-path lsp-name]**
- **[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name]**

You can enable the optimization timer for point-to-multipoint LSPs. See [“Optimizing Signaled LSPs” on page 251](#) for more information.

Configuring a Branch LSP for Point-to-Multipoint LSPs

The primary point-to-multipoint LSP sends traffic to two or more branch LSPs carrying traffic to each of the egress provider edge (PE) routers. In the configuration for each of these branch LSPs, the point-to-multipoint LSP path name you specify must be identical to the path name configured for the primary point-to-multipoint LSP. See [“Configuring the Primary Point-to-Multipoint LSP” on page 283](#) for more information.

To associate a branch LSP with the primary point-to-multipoint LSP, specify the point-to-multipoint LSP name by including the **p2mp** statement:

```
p2mp p2mp-lsp-name;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]



NOTE: Any change in any of the branch LSPs of a point-to-multipoint LSP, either due to a user action or an automatic adjustment made by the router, causes the primary and branch LSPs to be resigaled. The new point-to-multipoint LSP is signaled first before the old path is taken down.

The following sections describe how you can configure the branch LSP as a dynamically signaled path using Constrained Shortest Path First (CSPF), as a static path, or as a combination of dynamic and static paths:

- [Configuring the Branch LSP as a Dynamic Path on page 284](#)
- [Configuring the Branch LSP as a Static Path on page 284](#)

Configuring the Branch LSP as a Dynamic Path

By default, the branch LSP for a point-to-multipoint LSP is signaled dynamically using CSPF and requires no configuration.

When a point-to-multipoint LSP is changed, either by the addition or deletion of new destinations or by the recalculation of the path to existing destinations, certain nodes in the tree might receive data from more than one incoming interface. This can happen under the following conditions:

- Some of the branch LSPs to destinations are statically configured and might intersect with statically or dynamically calculated paths to other destinations.
- When a dynamically calculated path for a branch LSP results in a change of incoming interface for one of the nodes in the network, the older path is not immediately torn down after the new one has been signaled. This ensures that any data in transit relying on the older path can reach its destination. However, network traffic can potentially use either path to reach the destination.
- A faulty router at the ingress calculates the paths to two different branch destinations such that a different incoming interface is chosen for these branch LSPs on a router node common to these branch LSPs.

Configuring the Branch LSP as a Static Path

You can configure the branch LSP for a point-to-multipoint LSP as a static path. See [“Configuring Static LSPs” on page 271](#) for more information.

Example: Configuring a Collection of Paths to Create an RSVP-Signaled Point-to-Multipoint LSP

This example shows how to configure a collection of paths to create an RSVP-signaled point-to-multipoint label-switched path (LSP).

- [Requirements on page 285](#)
- [Overview on page 285](#)
- [Configuration on page 286](#)
- [Verification on page 302](#)

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

In this example, multiple routing devices serve as the transit, branch, and leaf nodes of a single point-to-multipoint LSP. On the provider edge (PE), Device PE1 is the ingress node. The branches go from PE1 to PE2, PE1 to PE3, and PE1 to PE4. Static unicast routes on the ingress node (PE1) point to the egress nodes.

This example also demonstrates static routes with a next hop that is a point-to-multipoint LSP, using the `p2mp-lsp-next-hop` statement. This is useful when implementing filter-based forwarding.




NOTE: Another option is to use the `lsp-next-hop` statement to configure a regular point-to-point LSP to be the next hop. Though not shown in this example, you can optionally assign an independent preference and metric to the next hop.

Topology Diagram

Figure 33 on page 286 shows the topology used in this example.

Figure 33: RSVP-Signaled Point-to-Multipoint LSP



g041174

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device PE1

```

set interfaces ge-2/0/2 unit 0 description PE1-to-CE1
set interfaces ge-2/0/2 unit 0 family inet address 10.0.244.10/30
set interfaces fe-2/0/10 unit 1 description PE1-to-P2
set interfaces fe-2/0/10 unit 1 family inet address 2.2.2.1/24
set interfaces fe-2/0/10 unit 1 family mpls
set interfaces fe-2/0/9 unit 8 description PE1-to-P3
set interfaces fe-2/0/9 unit 8 family inet address 6.6.6.1/24
set interfaces fe-2/0/9 unit 8 family mpls
set interfaces fe-2/0/8 unit 9 description PE1-to-P4
set interfaces fe-2/0/8 unit 9 family inet address 3.3.3.1/24
set interfaces fe-2/0/8 unit 9 family mpls
set interfaces lo0 unit 1 family inet address 100.10.10.10/32
set protocols rsvp interface fe-2/0/10.1
set protocols rsvp interface fe-2/0/9.8
set protocols rsvp interface fe-2/0/8.9
set protocols rsvp interface lo0.1
set protocols mpls traffic-engineering bgp-igp
set protocols mpls label-switched-path PE1-PE2 to 100.50.50.50
set protocols mpls label-switched-path PE1-PE2 link-protection
set protocols mpls label-switched-path PE1-PE2 p2mp p2mp1
set protocols mpls label-switched-path PE1-PE3 to 100.70.70.70
set protocols mpls label-switched-path PE1-PE3 link-protection
set protocols mpls label-switched-path PE1-PE3 p2mp p2mp1
set protocols mpls label-switched-path PE1-PE4 to 100.40.40.40
set protocols mpls label-switched-path PE1-PE4 link-protection

```

```

set protocols mpls label-switched-path PE1-PE4 p2mp p2mp1
set protocols mpls interface fe-2/0/10.1
set protocols mpls interface fe-2/0/9.8
set protocols mpls interface fe-2/0/8.9
set protocols mpls interface lo0.1
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-2/0/2.0
set protocols ospf area 0.0.0.0 interface fe-2/0/10.1
set protocols ospf area 0.0.0.0 interface fe-2/0/9.8
set protocols ospf area 0.0.0.0 interface fe-2/0/8.9
set protocols ospf area 0.0.0.0 interface lo0.1
set routing-options static route 5.5.5.0/24 p2mp-lsp-next-hop p2mp1
set routing-options static route 7.7.7.0/24 p2mp-lsp-next-hop p2mp1
set routing-options static route 4.4.4.0/24 p2mp-lsp-next-hop p2mp1
set routing-options router-id 100.10.10.10

```

Device CE1	<pre> set interfaces ge-1/3/2 unit 0 family inet address 10.0.244.9/30 set interfaces ge-1/3/2 unit 0 description CE1-to-PE1 set routing-options static route 10.0.104.8/30 next-hop 10.0.244.10 set routing-options static route 10.0.134.8/30 next-hop 10.0.244.10 set routing-options static route 10.0.224.8/30 next-hop 10.0.244.10 </pre>
Device CE2	<pre> set interfaces ge-1/3/3 unit 0 family inet address 10.0.224.9/30 set interfaces ge-1/3/3 unit 0 description CE2-to-PE2 set routing-options static route 10.0.244.8/30 next-hop 10.0.224.10 </pre>
Device CE3	<pre> set interfaces ge-2/0/1 unit 0 family inet address 10.0.134.9/30 set interfaces ge-2/0/1 unit 0 description CE3-to-PE3 set routing-options static route 10.0.244.8/30 next-hop 10.0.134.10 </pre>
Device CE4	<pre> set interfaces ge-3/1/3 unit 0 family inet address 10.0.104.10/30 set interfaces ge-3/1/3 unit 0 description CE4-to-PE4 set routing-options static route 10.0.244.8/30 next-hop 10.0.104.9 </pre>

Configuring the Ingress Label-Switched Router (LSR) (Device PE1)

Step-by-Step Procedure

To configure Device PE1:

1. Configure the interfaces, interface encapsulation, and protocol families.

```

[edit interfaces]
user@PE1# set ge-2/0/2 unit 0 description PE1-to-CE1
user@PE1# set ge-2/0/2 unit 0 family inet address 10.0.244.10/30
user@PE1# set fe-2/0/10 unit 1 description PE1-to-P2
user@PE1# set fe-2/0/10 unit 1 family inet address 2.2.2.1/24
user@PE1# set fe-2/0/10 unit 1 family mpls
user@PE1# set fe-2/0/9 unit 8 description PE1-to-P3
user@PE1# set fe-2/0/9 unit 8 family inet address 6.6.6.1/24
user@PE1# set fe-2/0/9 unit 8 family mpls
user@PE1# set fe-2/0/8 unit 9 description PE1-to-P4
user@PE1# set fe-2/0/8 unit 9 family inet address 3.3.3.1/24
user@PE1# set fe-2/0/8 unit 9 family mpls
user@PE1# set lo0 unit 1 family inet address 100.10.10.10/32

```

2. Enable RSVP, MPLS, and OSPF on the interfaces.

```
[edit protocols]
user@PE1# set rsvp interface fe-2/0/10.1
user@PE1# set rsvp interface fe-2/0/9.8
user@PE1# set rsvp interface fe-2/0/8.9
user@PE1# set rsvp interface lo0.1
user@PE1# set mpls interface fe-2/0/10.1
user@PE1# set mpls interface fe-2/0/9.8
user@PE1# set mpls interface fe-2/0/8.9
user@PE1# set mpls interface lo0.1
user@PE1# set ospf area 0.0.0.0 interface ge-2/0/2.0
user@PE1# set ospf area 0.0.0.0 interface fe-2/0/10.1
user@PE1# set ospf area 0.0.0.0 interface fe-2/0/9.8
user@PE1# set ospf area 0.0.0.0 interface fe-2/0/8.9
user@PE1# set ospf area 0.0.0.0 interface lo0.1
```

3. Configure the MPLS point-to-multipoint LSPs.

```
[edit protocols]
user@PE1# set mpls label-switched-path PE1-PE2 to 100.50.50.50
user@PE1# set mpls label-switched-path PE1-PE2 p2mp p2mp1
user@PE1# set mpls label-switched-path PE1-PE3 to 100.70.70.70
user@PE1# set mpls label-switched-path PE1-PE3 p2mp p2mp1
user@PE1# set mpls label-switched-path PE1-PE4 to 100.40.40.40
user@PE1# set mpls label-switched-path PE1-PE4 p2mp p2mp1
```

4. (Optional) Enable link protection on the LSPs.

Link protection helps to ensure that traffic sent over a specific interface to a neighboring router can continue to reach the router if that interface fails.

```
[edit protocols]
user@PE1# set mpls label-switched-path PE1-PE2 link-protection
user@PE1# set mpls label-switched-path PE1-PE3 link-protection
user@PE1# set mpls label-switched-path PE1-PE4 link-protection
```

5. Enable MPLS to perform traffic engineering for OSPF.

```
[edit protocols]
user@PE1# set mpls traffic-engineering bgp-igp
```

This causes the ingress routes to be installed in the inet.0 routing table. By default, MPLS performs traffic engineering for BGP only. You need to enable MPLS traffic engineering on the ingress LSR only.

6. Enable traffic engineering for OSPF.

```
[edit protocols]
user@PE1# set ospf traffic-engineering
```

This causes the shortest-path first (SPF) algorithm to take into account the LSPs configured under MPLS.

7. Configure the router ID.

```
[edit routing-options]
user@PE1# set router-id 100.10.10.10
```

8. Configure static IP unicast routes with the point-to-multipoint LSP name as the next hop for each route.

```
[edit routing-options]
user@PE1# set static route 5.5.5.0/24 p2mp-lsp-next-hop p2mp1
user@PE1# set static route 7.7.7.0/24 p2mp-lsp-next-hop p2mp1
user@PE1# set static route 4.4.4.0/24 p2mp-lsp-next-hop p2mp1
```

9. If you are done configuring the device, commit the configuration.

```
[edit]
user@PE1# commit
```

Configuring the Transit and Egress LSRs (Devices P2, P3, P4, PE2, PE3, and PE4)

Step-by-Step Procedure

To configure the transit and egress LSRs:

1. Configure the interfaces, interface encapsulation, and protocol families.

```
[edit]
user@P2# set interfaces fe-2/0/10 unit 2 description P2-to-PE1
user@P2# set interfaces fe-2/0/10 unit 2 family inet address 2.2.2.2/24
user@P2# set interfaces fe-2/0/10 unit 2 family mpls
user@P2# set interfaces fe-2/0/9 unit 10 description P2-to-PE2
user@P2# set interfaces fe-2/0/9 unit 10 family inet address 5.5.5.1/24
user@P2# set interfaces fe-2/0/9 unit 10 family mpls
user@P2# set interfaces lo0 unit 2 family inet address 100.20.20.20/32

user@PE2# set interfaces ge-2/0/3 unit 0 description PE2-to-CE2
user@PE2# set interfaces ge-2/0/3 unit 0 family inet address 10.0.224.10/30
user@PE2# set interfaces fe-2/0/10 unit 5 description PE2-to-P2
user@PE2# set interfaces fe-2/0/10 unit 5 family inet address 5.5.5.2/24
user@PE2# set interfaces fe-2/0/10 unit 5 family mpls
user@PE2# set interfaces lo0 unit 5 family inet address 100.50.50.50/32

user@P3# set interfaces fe-2/0/10 unit 6 description P3-to-PE1
user@P3# set interfaces fe-2/0/10 unit 6 family inet address 6.6.6.2/24
user@P3# set interfaces fe-2/0/10 unit 6 family mpls
user@P3# set interfaces fe-2/0/9 unit 11 description P3-to-PE3
user@P3# set interfaces fe-2/0/9 unit 11 family inet address 7.7.7.1/24
user@P3# set interfaces fe-2/0/9 unit 11 family mpls
user@P3# set interfaces lo0 unit 6 family inet address 100.60.60.60/32

user@PE3# set interfaces ge-2/0/1 unit 0 description PE3-to-CE3
user@PE3# set interfaces ge-2/0/1 unit 0 family inet address 10.0.134.10/30
user@PE3# set interfaces fe-2/0/10 unit 7 description PE3-to-P3
user@PE3# set interfaces fe-2/0/10 unit 7 family inet address 7.7.7.2/24
user@PE3# set interfaces fe-2/0/10 unit 7 family mpls
user@PE3# set interfaces lo0 unit 7 family inet address 100.70.70.70/32

user@P4# set interfaces fe-2/0/10 unit 3 description P4-to-PE1
user@P4# set interfaces fe-2/0/10 unit 3 family inet address 3.3.3.2/24
user@P4# set interfaces fe-2/0/10 unit 3 family mpls
user@P4# set interfaces fe-2/0/9 unit 12 description P4-to-PE4
user@P4# set interfaces fe-2/0/9 unit 12 family inet address 4.4.4.1/24
```

```
user@P4# set interfaces fe-2/0/9 unit 12 family mpls
user@P4# set interfaces lo0 unit 3 family inet address 100.30.30.30/32
```

```
user@PE4# set interfaces ge-2/0/0 unit 0 description PE4-to-CE4
user@PE4# set interfaces ge-2/0/0 unit 0 family inet address 10.0.104.9/30
user@PE4# set interfaces fe-2/0/10 unit 4 description PE4-to-P4
user@PE4# set interfaces fe-2/0/10 unit 4 family inet address 4.4.4.2/24
user@PE4# set interfaces fe-2/0/10 unit 4 family mpls
user@PE4# set interfaces lo0 unit 4 family inet address 100.40.40.40/32
```

2. Enable RSVP, MPLS, and OSPF on the interfaces.

```
[edit]
user@P2# set protocols rsvp interface fe-2/0/10.2
user@P2# set protocols rsvp interface fe-2/0/9.10
user@P2# set protocols rsvp interface lo0.2
user@P2# set protocols mpls interface fe-2/0/10.2
user@P2# set protocols mpls interface fe-2/0/9.10
user@P2# set protocols mpls interface lo0.2
user@P2# set protocols ospf area 0.0.0.0 interface fe-2/0/10.2
user@P2# set protocols ospf area 0.0.0.0 interface fe-2/0/9.10
user@P2# set protocols ospf area 0.0.0.0 interface lo0.2

user@PE2# set protocols rsvp interface fe-2/0/10.5
user@PE2# set protocols rsvp interface lo0.5
user@PE2# set protocols mpls interface fe-2/0/10.5
user@PE2# set protocols mpls interface lo0.5
user@PE2# set protocols ospf area 0.0.0.0 interface ge-2/0/3.0
user@PE2# set protocols ospf area 0.0.0.0 interface fe-2/0/10.5
user@PE2# set protocols ospf area 0.0.0.0 interface lo0.5

user@P3# set protocols rsvp interface fe-2/0/10.6
user@P3# set protocols rsvp interface fe-2/0/9.11
user@P3# set protocols rsvp interface lo0.6
user@P3# set protocols mpls interface fe-2/0/10.6
user@P3# set protocols mpls interface fe-2/0/9.11
user@P3# set protocols mpls interface lo0.6
user@P3# set protocols ospf area 0.0.0.0 interface fe-2/0/10.6
user@P3# set protocols ospf area 0.0.0.0 interface fe-2/0/9.11
user@P3# set protocols ospf area 0.0.0.0 interface lo0.6

user@PE3# set protocols rsvp interface fe-2/0/10.7
user@PE3# set protocols rsvp interface lo0.7
user@PE3# set protocols mpls interface fe-2/0/10.7
user@PE3# set protocols mpls interface lo0.7
user@PE3# set protocols ospf area 0.0.0.0 interface ge-2/0/1.0
user@PE3# set protocols ospf area 0.0.0.0 interface fe-2/0/10.7
user@PE3# set protocols ospf area 0.0.0.0 interface lo0.7

user@P4# set protocols rsvp interface fe-2/0/10.3
user@P4# set protocols rsvp interface fe-2/0/9.12
user@P4# set protocols rsvp interface lo0.3
user@P4# set protocols mpls interface fe-2/0/10.3
user@P4# set protocols mpls interface fe-2/0/9.12
```



```

user@P4# set protocols mpls interface lo0.3
user@P4# set protocols ospf area 0.0.0.0 interface fe-2/0/10.3
user@P4# set protocols ospf area 0.0.0.0 interface fe-2/0/9.12
user@P4# set protocols ospf area 0.0.0.0 interface lo0.3

```

```

user@PE4# set protocols rsvp interface fe-2/0/10.4
user@PE4# set protocols rsvp interface lo0.4
user@PE4# set protocols mpls interface fe-2/0/10.4
user@PE4# set protocols mpls interface lo0.4
user@PE4# set protocols ospf area 0.0.0.0 interface ge-2/0/0.0
user@PE4# set protocols ospf area 0.0.0.0 interface fe-2/0/10.4
user@PE4# set protocols ospf area 0.0.0.0 interface lo0.4

```

3. Enable traffic engineering for OSPF.

```

[edit]
user@P2# set protocols ospf traffic-engineering

```

```

user@P3# set protocols ospf traffic-engineering

```

```

user@P4# set protocols ospf traffic-engineering

```

```

user@PE2# set protocols ospf traffic-engineering

```

```

user@PE3# set protocols ospf traffic-engineering

```

```

user@PE4# set protocols ospf traffic-engineering

```

This causes the shortest-path first (SPF) algorithm to take into account the LSPs configured under MPLS.

4. Configure the router IDs.

```

[edit]
user@P2# set routing-options router-id 100.20.20.20

```

```

user@P3# set routing-options router-id 100.60.60.60

```

```

user@P4# set routing-options router-id 100.30.30.30

```

```

user@PE2# set routing-options router-id 100.50.50.50

```

```

user@PE3# set routing-options router-id 100.70.70.70

```

```

user@PE4# set routing-options router-id 100.40.40.40

```

5. If you are done configuring the devices, commit the configuration.

```

[edit]
user@host# commit

```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device PE1 user@PE1# show interfaces
ge-2/0/2 {
  unit 0 {
    description R1-to-CE1;
    family inet {
      address 10.0.244.10/30;
    }
  }
}
fe-2/0/10 {
  unit 1 {
    description PE1-to-P2;
    family inet {
      address 2.2.2.1/24;
    }
    family mpls;
  }
}
fe-2/0/9 {
  unit 8 {
    description PE1-to-P2;
    family inet {
      address 6.6.6.1/24;
    }
    family mpls;
  }
}
fe-2/0/8 {
  unit 9 {
    description PE1-to-P3;
    family inet {
      address 3.3.3.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 1 {
    family inet {
      address 100.10.10.10/32;
    }
  }
}

user@PE1# show protocols
rsvp {
  interface fe-2/0/10.1;
  interface fe-2/0/9.8;
  interface fe-2/0/8.9;
  interface lo0.1;
}
```

```

mpls {
  traffic-engineering bgp-igp;
  label-switched-path PE1-to-PE2 {
    to 100.50.50.50;
    link-protection;
    p2mp p2mp1;
  }
  label-switched-path PE1-to-PE3 {
    to 100.70.70.70;
    link-protection;
    p2mp p2mp1;
  }
  label-switched-path PE1-to-PE4 {
    to 100.40.40.40;
    link-protection;
    p2mp p2mp1;
  }
  interface fe-2/0/10.1;
  interface fe-2/0/9.8;
  interface fe-2/0/8.9;
  interface lo0.1;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-2/0/2.0;
    interface fe-2/0/10.1;
    interface fe-2/0/9.8;
    interface fe-2/0/8.9;
    interface lo0.1;
  }
}

user@PE1# show routing-options
static {
  route 5.5.5.0/24 {
    p2mp-lsp-next-hop p2mp1;
  }
  route 7.7.7.0/24 {
    p2mp-lsp-next-hop p2mp1;
  }
  route 4.4.4.0/24 {
    p2mp-lsp-next-hop p2mp1;
  }
}
router-id 100.10.10.10;

Device P2 user@P2# show interfaces
fe-2/0/10 {
  unit 2 {
    description P2-to-PE1;
    family inet {
      address 2.2.2.2/24;
    }
    family mpls;
  }
}
fe-2/0/9 {

```

```
    unit 10 {
      description P2-to-PE2;
      family inet {
        address 5.5.5.1/24;
      }
      family mpls;
    }
  }
  lo0 {
    unit 2 {
      family inet {
        address 100.20.20.20/32;
      }
    }
  }
}
```

user@P2# show protocols

```
rsvp {
  interface fe-2/0/10.2;
  interface fe-2/0/9.10;
  interface lo0.2;
}
mpls {
  interface fe-2/0/10.2;
  interface fe-2/0/9.10;
  interface lo0.2;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface fe-2/0/10.2;
    interface fe-2/0/9.10;
    interface lo0.2;
  }
}
```

user@P2# show routing-options
router-id 100.20.20.20;

Device P3

user@P3# show interfaces

```
fe-2/0/10 {
  unit 6 {
    description P3-to-PE1;
    family inet {
      address 6.6.6.2/24;
    }
    family mpls;
  }
}
fe-2/0/9 {
  unit 11 {
    description P3-to-PE3;
    family inet {
      address 7.7.7.1/24;
    }
    family mpls;
  }
}
```

```

    }
  }
  lo0 {
    unit 6 {
      family inet {
        address 100.60.60.60/32;
      }
    }
  }
}

```

user@P3# show protocols

```

rsvp {
  interface fe-2/0/10.6;
  interface fe-2/0/9.11;
  interface lo0.6;
}
mpls {
  interface fe-2/0/10.6;
  interface fe-2/0/9.11;
  interface lo0.6;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface fe-2/0/10.6;
    interface fe-2/0/9.11;
    interface lo0.6;
  }
}

```

user@P2# show routing-options
router-id 100.60.60.60;

Device P4

user@P4# show interfaces

```

fe-2/0/10 {
  unit 3 {
    description P4-to-PE1;
    family inet {
      address 3.3.3.2/24;
    }
    family mpls;
  }
}
fe-2/0/9 {
  unit 12 {
    description P4-to-PE4;
    family inet {
      address 4.4.4.1/24;
    }
    family mpls;
  }
}
lo0 {
  unit 3 {
    family inet {
      address 100.30.30.30/32;
    }
  }
}

```

```

    }
  }
}

user@P4# show protocols
rsvp {
  interface fe-2/0/10.3;
  interface fe-2/0/9.12;
  interface lo0.3;
}
mpls {
  interface fe-2/0/10.3;
  interface fe-2/0/9.12;
  interface lo0.3;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface fe-2/0/10.3;
    interface fe-2/0/9.12;
    interface lo0.3;
  }
}

```

```

user@P3# show routing-options
router-id 100.30.30.30;

```

Device PE2

```

user@PE2# show interfaces
ge-2/0/3 {
  unit 0 {
    description PE2-to-CE2;
    family inet {
      address 10.0.224.10/30;
    }
  }
}
fe-2/0/10 {
  unit 5 {
    description PE2-to-P2;
    family inet {
      address 5.5.5.2/24;
    }
    family mpls;
  }
}
lo0 {
  unit 5 {
    family inet {
      address 100.50.50.50/32;
    }
  }
}
}

user@PE2# show protocols
rsvp {

```

```

        interface fe-2/0/10.5;
        interface lo0.5;
    }
    mpls {
        interface fe-2/0/10.5;
        interface lo0.5;
    }
    ospf {
        traffic-engineering;
        area 0.0.0.0 {
            interface ge-2/0/3.0;
            interface fe-2/0/10.5;
            interface lo0.5;
        }
    }
}

user@PE2# show routing-options
router-id 100.50.50.50;

Device PE3 user@PE3# show interfaces
ge-2/0/1 {
    unit 0 {
        description PE3-to-CE3;
        family inet {
            address 10.0.134.10/30;
        }
    }
}
fe-2/0/10 {
    unit 7 {
        description PE3-to-P3;
        family inet {
            address 7.7.7.2/24;
        }
        family mpls;
    }
}
lo0 {
    unit 7 {
        family inet {
            address 100.70.70.70/32;
        }
    }
}
}

user@PE3# show protocols
rsvp {
    interface fe-2/0/10.7;
    interface lo0.7;
}
mpls {
    interface fe-2/0/10.7;
    interface lo0.7;
}
ospf {

```

```

        traffic-engineering;
        area 0.0.0.0 {
            interface ge-2/0/1.0;
            interface fe-2/0/10.7;
            interface lo0.7;
        }
    }

user@PE3# show routing-options
router-id 100.70.70.70;

Device PE4 user@PE4# show interfaces
ge-2/0/0 {
    unit 0 {
        description PE4-to-CE4;
        family inet {
            address 10.0.104.9/30;
        }
    }
}
fe-2/0/10 {
    unit 4 {
        description PE4-to-P4;
        family inet {
            address 4.4.4.2/24;
        }
        family mpls;
    }
}
lo0 {
    unit 4 {
        family inet {
            address 100.40.40.40/32;
        }
    }
}

user@PE4# show protocols
rsvp {
    interface fe-2/0/10.4;
    interface lo0.4;
}
mpls {
    interface fe-2/0/10.4;
    interface lo0.4;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface ge-2/0/0.0;
        interface fe-2/0/10.4;
        interface lo0.4;
    }
}

user@PE4# show routing-options

```



```
router-id 100.40.40.40;
```

Configuring Device CE1

Step-by-Step Procedure

To configure Device CE1:

1. Configure an interface to Device PE1.

```
[edit interfaces]
user@CE1# set ge-1/3/2 unit 0 family inet address 10.0.244.9/30
user@CE1# set ge-1/3/2 unit 0 description CE1-to-PE1
```

2. Configure static routes from Device CE1 to the three other customer networks, with Device PE1 as the next hop.

```
[edit routing-options]
user@CE1# set static route 10.0.104.8/30 next-hop 10.0.244.10
user@CE1# set static route 10.0.134.8/30 next-hop 10.0.244.10
user@CE1# set static route 10.0.224.8/30 next-hop 10.0.244.10
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE1# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE1# show interfaces
ge-1/3/2 {
  unit 0 {
    family inet {
      address 10.0.244.9/30;
      description CE1-to-PE1;
    }
  }
}

user@CE1# show routing-options
static {
  route 10.0.104.8/30 next-hop 10.0.244.10;
  route 10.0.134.8/30 next-hop 10.0.244.10;
  route 10.0.224.8/30 next-hop 10.0.244.10;
}
```

Configuring Device CE2

Step-by-Step Procedure

To configure Device CE2:

1. Configure an interface to Device PE2.

```
[edit interfaces]
user@CE2# set ge-1/3/3 unit 0 family inet address 10.0.224.9/30
user@CE2# set ge-1/3/3 unit 0 description CE2-to-PE2
```

2. Configure a static route from Device CE2 to CE1, with Device PE2 as the next hop.

```
[edit routing-options]
user@CE2# set static route 10.0.244.8/30 next-hop 10.0.224.10
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE2# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE2# show interfaces
ge-1/3/3 {
  unit 0 {
    family inet {
      address 10.0.224.9/30;
      description CE2-to-PE2;
    }
  }
}

user@CE2# show routing-options
static {
  route 10.0.244.8/30 next-hop 10.0.224.10;
}
```

Configuring Device CE3

Step-by-Step Procedure To configure Device CE3:

1. Configure an interface to Device PE3.

```
[edit interfaces]
user@CE3# set ge-2/0/1 unit 0 family inet address 10.0.134.9/30
user@CE3# set ge-2/0/1 unit 0 description CE3-to-PE3
```

2. Configure a static route from Device CE3 to CE1, with Device PE3 as the next hop.

```
[edit routing-options]
user@CE3# set static route 10.0.244.8/30 next-hop 10.0.134.10
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@CE3# commit
```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE3# show interfaces
ge-2/0/1 {
  unit 0 {
```

```

        family inet {
            address 10.0.134.9/30;
            description CE3-to-PE3;
        }
    }
}

user@CE3# show routing-options
static {
    route 10.0.244.8/30 next-hop 10.0.134.10;
}

```

Configuring Device CE4

Step-by-Step Procedure

To configure Device CE4:

1. Configure an interface to Device PE4.

```

[edit interfaces]
user@CE4# set ge-3/1/3 unit 0 family inet address 10.0.104.10/30
user@CE4# set ge-3/1/3 unit 0 description CE4-to-PE4

```

2. Configure a static route from Device CE4 to CE1, with Device PE4 as the next hop.

```

[edit routing-options]
user@CE4# set static route 10.0.244.8/30 next-hop 10.0.104.9

```

3. If you are done configuring the device, commit the configuration.

```

[edit]
user@CE4# commit

```

Results From configuration mode, confirm your configuration by entering the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@CE4# show interfaces
ge-3/1/3 {
    unit 0 {
        family inet {
            address 10.0.104.10/30;
            description CE4-to-PE4;
        }
    }
}

user@CE4# show routing-options
static {
    route 10.0.244.8/30 next-hop 10.0.104.9;
}

```

Verification

Confirm that the configuration is working properly.

- [Verifying Connectivity on page 302](#)
- [Verifying the State of the Point-to-Multipoint LSP on page 302](#)
- [Checking the Forwarding Table on page 303](#)

Verifying Connectivity

Purpose Make sure that the devices can ping each other.

Action Run the **ping** command from CE1 to the interface on CE2 connecting to PE2.

```
user@CE1> ping 10.0.224.9
PING 10.0.224.9 (10.0.224.9): 56 data bytes
64 bytes from 10.0.224.9: icmp_seq=0 ttl=61 time=1.387 ms
64 bytes from 10.0.224.9: icmp_seq=1 ttl=61 time=1.394 ms
64 bytes from 10.0.224.9: icmp_seq=2 ttl=61 time=1.506 ms
^C
--- 10.0.224.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.387/1.429/1.506/0.055 ms
```

Run the **ping** command from CE1 to the interface on CE3 connecting to PE3.

```
user@CE1> ping 10.0.134.9
PING 10.0.134.9 (10.0.134.9): 56 data bytes
64 bytes from 10.0.134.9: icmp_seq=0 ttl=61 time=1.068 ms
64 bytes from 10.0.134.9: icmp_seq=1 ttl=61 time=1.062 ms
64 bytes from 10.0.134.9: icmp_seq=2 ttl=61 time=1.053 ms
^C
--- 10.0.134.9 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.053/1.061/1.068/0.006 ms
```

Run the **ping** command from CE1 to the interface on CE4 connecting to PE4.

```
user@CE1> ping 10.0.104.10
PING 10.0.104.10 (10.0.104.10): 56 data bytes
64 bytes from 10.0.104.10: icmp_seq=0 ttl=61 time=1.079 ms
64 bytes from 10.0.104.10: icmp_seq=1 ttl=61 time=1.048 ms
64 bytes from 10.0.104.10: icmp_seq=2 ttl=61 time=1.070 ms
^C
--- 10.0.104.10 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.048/1.066/1.079/0.013 ms
```

Verifying the State of the Point-to-Multipoint LSP

Purpose Make sure that the ingress, transit, and egress LSRs are in the Up state.

Action Run the `show mpls lsp p2mp` command on all of the LSRs. Only the ingress LSR is shown here.

```
user@PE1> show mpls lsp p2mp
Ingress LSP: 1 sessions
P2MP name: p2mp1, P2MP branch count: 3
To          From          State Rt P    ActivePath    LSPname
100.40.40.40 100.10.10.10 Up    0 *          PE1-PE4
100.70.70.70 100.10.10.10 Up    0 *          PE1-PE3
100.50.50.50 100.10.10.10 Up    0 *          PE1-PE2
Total 3 displayed, Up 3, Down 0
...
```

Checking the Forwarding Table

Purpose Make sure that the routes are set up as expected by running the `show route forwarding-table` command. Only the routes to the remote customer networks are shown here.

Action user@PE1> show route forwarding-table

```
Routing table: default.inet
Internet:
Destination          Type RtRef Next hop          Type Index NhRef Netif
...
10.0.104.8/30         user    0 3.3.3.2          ucst  1006   6 fe-2/0/8.9
10.0.134.8/30         user    0 6.6.6.2          ucst  1010   6 fe-2/0/9.8
10.0.224.8/30         user    0 2.2.2.2          ucst  1008   6 fe-2/0/10.1
...
```

Related Documentation • [Point-to-Multipoint LSPs Overview on page 281](#) in the *Junos OS MPLS Applications Library for Routing Devices*

Configuring Inter-Domain Point-to-Multipoint LSPs

An inter-domain P2MP LSP is a P2MP LSP that has one or more sub-LSPs (branches) that span multiple domains in a network. Examples of such domains include IGP areas and autonomous systems (ASs). A sub-LSP of an inter-domain P2MP LSP may be intra-area, inter-area, or inter-AS, depending on the location of the egress node (leaf) with respect to the ingress node (source).

On the ingress node, a name is assigned to the inter-domain P2MP LSP and shared by all constituent sub-LSPs. Each sub-LSP is configured separately, with its own egress node and optionally an explicit path. The location of the egress node of the sub-LSP with respect to the ingress node determines whether the sub-LSP is intra-area, inter-area, or inter-AS.

Inter-domain P2MP LSPs can be used to transport traffic in the following applications in a multi-area or multi-AS network:

- Layer 2 broadcast and multicast over MPLS
- Layer 3 BGP/MPLS VPN

- VPLS

On each domain boundary node (ABR or ASBR) along the path of the P2MP LSP, the **expand-loose-hop** statement must be configured at the **[edit protocols mpls]** hierarchy level so that CSPF can extend a loose-hop ERO (usually the first entry of the ERO list carried by RSVP Path message) towards the egress node or the next domain boundary node.

CSPF path computation for inter-domain P2MP LSPs:

- CSPF path computation is supported on each sub-LSP for inter-domain P2MP LSPs. A sub-LSP may be intra-area, inter-area, or inter-AS. CSPF treats an inter-area or inter-AS sub-LSP in the same manner as an inter-domain P2P LSP.
- On an ingress node or a domain boundary node (ABR or ASBR), CSPF can perform an Explicit Route Object (ERO) expansion per-RSVP query. The destination queried could be an egress node or a received loose-hop ERO. If the destination resides in a neighboring domain that the node is connected to, CSPF generates either a sequence of strict-hop EROs towards it or a sequence of strict-hop EROs towards another domain boundary node that can reach the destination.
- If RSVP fails to signal a path through a previously selected domain boundary node, RSVP attempts to signal a path through other available domain boundary nodes in a round-robin fashion.
- When a sub-LSP is added or removed to or from an inter-domain P2MP LSP, causing its path (branch) to be merged or pruned with or from the current P2MP tree, the paths being taken by the other sub-LSPs should not be affected, helping to prevent traffic disruption on those sub-LSPs.

Be aware of the following when deploying inter-domain P2MP LSPs in your network:

- Periodic path re-optimization is supported for inter-domain P2MP LSPs on ingress nodes. It can be turned on for an inter-domain P2MP LSP by configuring the **optimize-timer** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level with the same interval for every sub-LSP.
- Only link protection bypass LSPs are supported for inter-domain P2MP LSPs. To enable it for an inter-domain P2MP LSP, link-protection must be configured for all sub-LSPs and on all of the RSVP interfaces that the P2MP LSP might travel through.
- Only OSPF areas are supported for inter-domain P2MP LSPs. IS-IS levels are not supported.

Configuring Link Protection for Point-to-Multipoint LSPs

Link protection helps to ensure that traffic going over a specific interface to a neighboring router can continue to reach this router if that interface fails. When link protection is configured for an interface and a point-to-multipoint LSP that traverses this interface, a bypass LSP is created that handles this traffic if the interface fails. The bypass LSP uses a different interface and path to reach the same destination.

To extend link protection to all of the paths used by a point-to-multipoint LSP, link protection must be configured on each router that each branch LSP traverses. If you enable link protection on a point-to-multipoint LSP, you must enable link protection on all of the branch LSPs.

The Internet draft [draft-ietf-mpls-rsvp-te-p2mp-01.txt](#), *Extensions to RSVP-TE for Point to Multipoint TE LSPs*, describes link protection for point-to-multipoint LSPs.

To enable link protection on point-to-multipoint LSPs, complete the following steps:

1. Configure link protection on each branch LSP. To configure link protection, include the **link-protection** statement:

```
link-protection;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls [label-switched-path](#) *branch-lsp-name*]
 - [edit logical-systems *logical-system-name* protocols mpls [label-switched-path](#) *branch-lsp-name*]
2. Configure link protection for each RSVP interface on each router that the branch LSP traverses. For information about how to configure link protection on RSVP interfaces, see [“Configuring Link Protection on Interfaces Used by LSPs” on page 502](#).

For more information on how to configure link protection, see [“Configuring Node Protection or Link Protection for LSPs” on page 509](#).

Configuring Graceful Restart for Point-to-Multipoint LSPs

You can configure graceful restart on point-to-multipoint LSPs. Graceful restart allows a router undergoing a restart to inform its adjacent neighbors of its condition. The restarting router requests a grace period from the neighbor or peer, which can then cooperate with the restarting router. The restarting router can still forward MPLS traffic during the restart period; convergence in the network is not disrupted. The restart is not apparent to the rest of the network, and the restarting router is not removed from the network topology. RSVP graceful restart can be enabled on both transit routers and ingress routers.

To enable graceful restart on a router handling point-to-multipoint LSP traffic, include the **graceful-restart** statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- [edit routing-options]
- [edit logical-systems *logical-system-name* routing-options]

The graceful restart configuration for point-to-multipoint LSPs is identical to that of point-to-point LSPs. For more information on how to configure graceful restart, see [“Configuring RSVP Graceful Restart” on page 514](#).

Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs

You can control whether a reverse path forwarding (RPF) check is performed for a source and group entry before installing a route in the multicast forwarding cache. This makes it possible to use point-to-multipoint LSPs to distribute multicast traffic to PIM islands situated downstream from the egress routers of the point-to-multipoint LSPs.

By configuring the **rpf-check-policy** statement, you can disable RPF checks for a source and group pair. You would typically configure this statement on the egress routers of a point-to-multipoint LSP, because the interface receiving the multicast traffic on a point-to-multipoint LSP egress router might not always be the RPF interface.

You can also configure a routing policy to act upon a source and group pair. This policy behaves like an import policy, so if no policy term matches the input data, the default policy action is “acceptance.” An accept policy action enables RPF checks. A reject policy action (applied to all source and group pairs that are not accepted) disables RPF checks for the pair.

To configure a multicast RPF check policy for a point-to-multipoint LSP, specify the RPF check policy using the **rpf-check-policy** statement:

```
rpf-check-policy policy;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options multicast]**
- **[edit logical-systems *logical-system-name* routing-options multicast]**

You also must configure a policy for the multicast RPF check. You configure policies at the **[edit policy-options]** hierarchy level. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.



NOTE: When you configure the **rpf-check-policy** statement, the Junos OS cannot perform RPF checks on incoming traffic and therefore cannot detect traffic arriving on the wrong interface. This might cause routing loops to form.

Example: Configuring Multicast RPF Check Policy for a Point-to-Multipoint LSP

Configure a policy to ensure that an RPF check is not performed for sources with prefix **128.83/16** or longer that belong to groups having a prefix of **228/8** or longer:

```
[edit]
policy-options {
  policy-statement rpf-sg-policy {
    from {
      route-filter 228.0.0.0/8 orlonger;
      source-address-filter 128.83.0.0/16 orlonger;
    }
    then {
      reject;
    }
  }
}
```



```

    }
  }
}

```

Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs

You can configure one or more PE routers as part of a backup PE router group to enable ingress PE router redundancy. You accomplish this by configuring the IP addresses of the backup PE routers (at least one backup PE router is required) and the local IP address used by the local PE router.

You must also configure a full mesh of point-to-point LSPs between the primary and backup PE routers. You also need to configure BFD on these LSPs. See [“Configuring BFD for RSVP-Signaled LSPs” on page 355](#) and [“Configuring BFD for LDP LSPs” on page 543](#) for more information.

To configure ingress PE router redundancy for point-to-multipoint LSPs, include the **backup-pe-group** statement:

```

backup-pe-group pe-group-name {
  backups [addresses];
  local-address address;
}

```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

After you configure the ingress PE router redundancy backup group, you must also apply the group to a static route on the PE router. This ensures that the static route is active (installed in the forwarding table) when the local PE router is the designated forwarder for the backup PE group. You can only associate a backup PE router group with a static route that also has the **p2mp-lsp-next-hop** statement configured. For more information, see [“Configuring Static Unicast Routes for Point-to-Multipoint LSPs” on page 277](#).

Enabling Point-to-Point LSPs to Monitor Egress PE Routers

Configuring an LSP with the **associate-backup-pe-groups** statement enables it to monitor the status of the PE router to which it is configured. You can configure multiple backup PE router groups using the same router's address. A failure of this LSP indicates to all of the backup PE router groups that the destination PE router is down. The **associate-backup-pe-groups** statement is not tied to a specific backup PE router group. It applies to all groups that are interested in the status of the LSP to that address.

To allow an LSP to monitor the status of the egress PE router, include the **associate-backup-pe-groups** statement:

```

associate-backup-pe-groups;

```

This statement can be configured at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

If you configure the **associate-backup-pe-groups** statement, you must configure BFD for the point-to-point LSP. For information about how to configure BFD for an LSP, see [“Configuring BFD for MPLS IPv4 LSPs” on page 354](#) and [“Configuring BFD for LDP LSPs” on page 543](#).

You also must configure a full mesh of point-to-point LSPs between the PE routers in the backup PE router group. A full mesh is required so that each PE router within the group can independently determine the status of the other PE routers, allowing each router to independently determine which PE router is currently the designated forwarder for the backup PE router group.

If you configure multiple LSPs with the **associate-backup-pe-groups** statement to the same destination PE router, the first LSP configured is used to monitor the forwarding state to that PE router. If you configure multiple LSPs to the same destination, make sure to configure similar parameters for the LSPs. With this configuration scenario, a failure notification might be triggered even though the remote PE router is still up.

Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases

In Junos OS Release 9.1 and earlier, Resv messages that include the S2L_SUB_LSP object are rejected by default. In Junos OS Release 9.2 and later, such messages are accepted by default. To ensure proper functioning of point-to-multipoint LSPs in a network that includes both devices running Junos OS Release 9.1 and earlier and devices running Junos 9.2 and later, you must include the **no-p2mp-sublsp** statement in the configuration of the devices running Junos 9.2 and later:

```
no-p2mp-sublsp;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols rsvp]**
- **[edit logical-systems *logical-system-name* protocols rsvp]**

CHAPTER 7

Configuring DiffServ-Aware Traffic Engineering to Achieve Service Level Guarantees on an MPLS network

- [DiffServ-Aware Traffic Engineering Introduction on page 310](#)
- [DiffServ-Aware Traffic Engineering Standards on page 310](#)
- [DiffServ-Aware Traffic Engineering Terminology on page 310](#)
- [DiffServ-Aware Traffic Engineering Features on page 311](#)
- [DiffServ-Aware Traffic Engineered LSPs on page 312](#)
- [DiffServ-Aware Traffic Engineered LSPs Overview on page 312](#)
- [DiffServ-Aware Traffic Engineered LSPs Operation on page 313](#)
- [Multiclass LSPs on page 313](#)
- [Multiclass LSP Overview on page 314](#)
- [Establishing a Multiclass LSP on the Differentiated Services Domain on page 314](#)
- [Configuring Routers for DiffServ-Aware Traffic Engineering on page 315](#)
- [LSP Bandwidth Oversubscription Overview on page 319](#)
- [LSP Size Oversubscription on page 320](#)
- [LSP Link Size Oversubscription on page 320](#)
- [Class Type Oversubscription and Local Oversubscription Multipliers on page 320](#)
- [Class Type Bandwidth and the LOM on page 321](#)
- [LOM Calculation for the MAM and Extended MAM Bandwidth Models on page 321](#)
- [LOM Calculation for the Russian Dolls Bandwidth Model on page 322](#)
- [Example: LOM Calculation on page 322](#)
- [Configuring the Bandwidth Subscription Percentage for LSPs on page 323](#)
- [Configuring LSPs for DiffServ-Aware Traffic Engineering on page 325](#)
- [Configuring Multiclass LSPs on page 328](#)

DiffServ-Aware Traffic Engineering Introduction

Differentiated Services (DiffServ)-aware traffic engineering provides a way to guarantee a specified level of service over an MPLS network. The routers providing DiffServ-aware traffic engineering are part of a differentiated services network domain. All routers participating in a differentiated services domain must have DiffServ-aware traffic engineering enabled.

To help ensure that the specified service level is provided, it is necessary to ensure that no more than the amount of traffic specified is sent over the differentiated services domain. You can accomplish this goal by configuring a policer to police or rate-limit the volume of traffic transiting the differentiated service domain. For more information about how to configure policers for label-switched paths (LSPs), see “[Configuring Policers for LSPs](#)” on page 347.

This feature can help to improve the quality of Internet services such as voice over IP (VoIP). It also makes it possible to better emulate an Asynchronous Transfer Mode (ATM) circuit over an MPLS network.

DiffServ-Aware Traffic Engineering Standards

The following RFCs provide information on DiffServ-aware traffic engineering and multiclass LSPs:

- RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*
- RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*
- RFC 4124, *Protocol Extensions for Support of Differentiated-Service-Aware MPLS Traffic Engineering*
- RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diff-Serv-aware MPLS Traffic Engineering*
- RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diff-Serv-aware MPLS*

These RFCs are available on the IETF website at <http://www.ietf.org/>.

DiffServ-Aware Traffic Engineering Terminology

B

Bandwidth model The bandwidth model determines the values of the available bandwidth advertised by the interior gateway protocols (IGPs).

C

CAC Call admission control (CAC) checks to ensure there is adequate bandwidth on the path before the LSP is established. If the bandwidth is insufficient, the LSP is not established and an error is reported.

Class type	A collection of traffic flows that is treated equivalently in a differentiated services domain. A class type maps to a queue and is much like a class-of-service (CoS) forwarding class in concept. It is also known as a traffic class.
D	
Differentiated Services	Differentiated Services make it possible to give different treatment to traffic based on the EXP bits in the MPLS header. Traffic must be marked appropriately and CoS must be configured.
Differentiated Services domain	The routers in a network that have Differentiated Services enabled.
DiffServ-aware traffic engineering	A type of constraint-based routing. It can enforce different bandwidth constraints for different classes of traffic. It can also do CAC on each traffic engineering class when an LSP is established.
M	
MAM	The maximum allocation bandwidth constraint model divides the available bandwidth between the different classes. Sharing of bandwidth between the class types is not allowed.
Multiclass LSP	A multiclass LSP functions like a standard LSP, but it also allows you to reserve bandwidth from multiple class types. The EXP bits of the MPLS header are used to distinguish between class types.
R	
RDM	The Russian dolls bandwidth constraint model makes efficient use of bandwidth by allowing the class types to share bandwidth.
T	
Traffic engineering class	A paired class type and priority.
Traffic engineering class map	A map between the class types, priorities, and traffic engineering classes. The traffic engineering class mapping must be consistent across the Differentiated Services domain.

DiffServ-Aware Traffic Engineering Features

DiffServ-aware traffic engineering provides the following features:

- Traffic engineering at a per-class level rather than at an aggregate level
- Different bandwidth constraints for different class types (traffic classes)
- Different queuing behaviors per class, allowing the router to forward traffic based on the class type

In comparison, standard traffic engineering does not consider CoS, and it completes its work on an aggregate basis across all Differentiated Service classes.

DiffServ-aware traffic engineering provides the following advantages:

- Traffic engineering can be performed on a specific class type instead of at the aggregate level.
- Bandwidth constraints can be enforced on each specific class type.
- It forwards traffic based on the EXP bits.

This makes it possible to guarantee service and bandwidth across an MPLS network. With DiffServ-aware traffic engineering, among other services, you can provide ATM circuit emulation, VoIP, and a guaranteed bandwidth service.

The following describes how the IGP, Constrained Shortest Path First (CSPF), and RSVP participate in DiffServ-aware traffic engineering:

- The IGP can advertise the unreserved bandwidth for each traffic engineering class to the other members of the differentiated services domain. The traffic engineering database stores this information.
- A CSPF calculation is performed considering the bandwidth constraints for each class type. If all the constraints are met, the CSPF calculation is considered successful.
- When RSVP signals an LSP, it requests bandwidth for specified class types.

DiffServ-Aware Traffic Engineered LSPs

A DiffServ-aware traffic engineered LSP is an LSP configured to reserve bandwidth for one of the supported class types and to carry traffic for that class type. The following sections discuss this type of LSPs:

- [DiffServ-Aware Traffic Engineered LSPs Overview on page 312](#)
- [DiffServ-Aware Traffic Engineered LSPs Operation on page 313](#)

DiffServ-Aware Traffic Engineered LSPs Overview

A DiffServ-aware traffic engineered LSP is an LSP configured with a bandwidth reservation for a specific class type. This LSP can carry traffic for a single class type. On the packets, the class type is specified by the EXP bits (also known as the class-of-service bits) and the per-hop behavior (PHB) associated with the EXP bits. The mapping between the EXP bits and the PHB is static, rather than being signaled in RSVP.

The class type must be configured consistently across the Differentiated Services domain, meaning the class type configuration must be consistent from router to router in the network. You can unambiguously map a class type to a queue. On each node router, the class-of-service queue configuration for an interface translates to the available bandwidth for a particular class type on that link.

For more information about topics related to LSPs and DiffServ-aware traffic engineering, see the following:

- For forwarding classes and class of service, see the *Class of Service Feature Guide for Routing Devices*.
- For EXP bits, see [“MPLS Label Allocation” on page 26](#).
- For differentiated services, see RFC 3270, *Multi-Protocol Label Switching (MPLS) Support of Differentiated Services*.
- For information about how the IGPs and RSVP have been modified to support Differentiated Services-aware MPLS traffic engineering, see RFC 4124, *Protocol Extensions for Support of Differentiated-Service-Aware MPLS Traffic Engineering*.

DiffServ-Aware Traffic Engineered LSPs Operation

When configuring a DiffServ-aware traffic engineered LSP, you specify the class type and the bandwidth associated with it. The following occurs when an LSP is established with bandwidth reservation from a specific class type:

1. The IGPs advertise how much unreserved bandwidth is available for the traffic engineering classes.
2. When calculating the path for an LSP, CSPF is used to ensure that the bandwidth constraints are met for the class type carried by the LSP at the specified priority level.

CSPF also checks to ensure that the bandwidth model is configured consistently on each router participating in the LSP. If the bandwidth model is inconsistent, CSPF does not compute the path (except for LSPs from class type ct0).
3. Once a path is found, RSVP signals the LSP using the Classtype object in the path message. At each node in the path, the available bandwidth for the class types is adjusted as the path is set up.

An LSP that requires bandwidth from a particular class (except class type ct0) cannot be established through routers that do not understand the Classtype object. Preventing the use of routers that do not understand the Classtype object helps to ensure consistency throughout the Differentiated Services domain by preventing the LSP from using a router that cannot support Differentiated Services.

By default, LSPs are signaled with setup priority 7 and holding priority 0. An LSP configured with these values cannot preempt another LSP at setup time and cannot be preempted.

It is possible to have both LSPs configured for DiffServ-aware traffic engineering and regular LSPs configured at the same time on the same physical interfaces. For this type of heterogeneous environment, regular LSPs carry best-effort traffic by default. Traffic carried in the regular LSPs must have the correct EXP settings (either by remarking the EXP settings or by assuming that the traffic arrived with the correct EXP settings from the upstream router).

Multiclass LSPs

Multiclass LSPs function like standard LSPs, but they also allow you to configure multiple class types with guaranteed bandwidth. The EXP bits of the MPLS header are used to distinguish between class types. Multiclass LSPs can be configured for a variety of

purposes. For example, you can configure a multiclass LSP to emulate the behavior of an ATM circuit. An ATM circuit can provide service-level guarantees to a class type. A multiclass LSP can provide a similar guaranteed level of service.

The following sections discuss multiclass LSPs:

- [Multiclass LSP Overview on page 314](#)
- [Establishing a Multiclass LSP on the Differentiated Services Domain on page 314](#)

Multiclass LSP Overview

A multiclass LSP is an LSP that can carry several class types. One multiclass LSP can be used to support up to four class types. On the packets, the class type is specified by the EXP bits (also known as the class-of-service bits) and the per-hop behavior (PHB) associated with the EXP bits. The mapping between the EXP bits and the PHB is static, rather than being signaled in RSVP.

Once a multiclass LSP is configured, traffic from all of the class types can:

- Follow the same path
- Be rerouted along the same path
- Be taken down at the same time

Class types must be configured consistently across the Differentiated Services domain, meaning the class type configuration must be consistent from router to router in the network.

You can unambiguously map a class type to a queue. On each node router, the CoS queue configuration for an interface translates to the available bandwidth for a particular class type on that link.

The combination of a class type and a priority level forms a traffic engineering class. The IGP can advertise up to eight traffic engineering classes for each link.

For more information about the EXP bits, see [“MPLS Label Allocation” on page 26](#).

For more information about forwarding classes, see the *Class of Service Feature Guide for Routing Devices*.

Establishing a Multiclass LSP on the Differentiated Services Domain

The following occurs when a multiclass LSP is established on the differentiated services domain:

1. The IGPs advertise how much unreserved bandwidth is available for the traffic engineering classes.
2. When calculating the path for a multiclass LSP, CSPF is used to ensure that the constraints are met for all the class types carried by the multiclass LSP (a set of constraints instead of a single constraint).

3. Once a path is found, RSVP signals the LSP using an RSVP object in the path message. At each node in the path, the available bandwidth for the class types is adjusted as the path is set up. The RSVP object is a hop-by-hop object. Multiclass LSPs cannot be established through routers that do not understand this object. Preventing routers that do not understand the RSVP object from carrying traffic helps to ensure consistency throughout the differentiated services domain by preventing the multiclass LSP from using a router that is incapable of supporting differentiated services.

By default, multiclass LSPs are signaled with setup priority 7 and holding priority 0. A multiclass LSP configured with these values cannot preempt another LSP at setup time and cannot be preempted.

It is possible to have both multiclass LSPs and regular LSPs configured at the same time on the same physical interfaces. For this type of heterogeneous environment, regular LSPs carry best-effort traffic by default. Traffic carried in the regular LSPs must have the correct EXP settings.

Configuring Routers for DiffServ-Aware Traffic Engineering

To configure DiffServ-aware traffic engineering, include the **diffserv-te** statement:

```
diffserv-te {
  bandwidth-model {
    extended-mam;
    mam;
    rdm;
  }
  te-class-matrix {
    traffic-class {
      tenumber {
        priority priority;
        traffic-class ctnumber priority priority;
      }
    }
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

You must include the **diffserv-te** statement in the configuration on all routers participating in the Differentiated Services domain. However, you are not required to configure the traffic engineering class matrix (by including the **te-class-matrix** statement at the [edit protocols mpls **diffserv-te**] or [edit logical-systems *logical-system-name* protocols mpls **diffserv-te**] hierarchy level).



NOTE: To prevent the possibility of an incorrect configuration when migrating to Diffserv-aware traffic engineering, a policy control failure error might be triggered if there is conflict between the old LSPs and the newly configured TE-class matrix.

An old node might request an LSP with setup and hold priorities in such a way that the combination of the ct0 class and the priority does not match with the configured TE-class matrix. All LSPs on the router that are configured prior to configuring diffserv-aware traffic engineering are designated as being from class ct0.

The error appears in the RSVP tracing logs as a **Session preempted** error. For the router where the error originates, the error could appear as follows:

```
Jun 17 16:35:59 RSVP error for session 10.255.245.6(port/tunnel ID 31133)
  Proto 0: (class ct0, priority 2) is not a valid TE-class Jun 17
16:35:59 RSVP originate PathErr 192.168.37.22->192.168.37.23 Session
preempted
```

For the router receiving the error, the error can appear as follows:

```
Jun 17 16:37:51 RSVP recv PathErr 192.168.37.22->192.168.37.23 Session
preempted LSP to-f(2/31133)
```

To configure DiffServ-aware traffic engineering, complete the procedures in the following sections:

- [Configuring the Bandwidth Model on page 316](#)
- [Configuring Traffic Engineering Classes on page 317](#)
- [Configuring Class of Service for DiffServ-Aware Traffic Engineering on page 318](#)

Configuring the Bandwidth Model

You must configure a bandwidth model on all routers participating in the Differentiated Services domain. The bandwidth models available are MAM, extended MAM, and RDM:

- Maximum allocation bandwidth constraints model (MAM)—Defined in RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*.
- Extended MAM—A proprietary bandwidth model that behaves much like standard MAM. If you configure multiclass LSPs, you must configure the extended MAM bandwidth model.
- Russian-dolls bandwidth allocation model (RDM)—Makes efficient use of bandwidth by allowing the class types to share bandwidth. RDM is defined in RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*.

To configure a bandwidth model, include the **bandwidth-model** statement and specify one of the bandwidth model options:

```
bandwidth-model {
  extended-mam;
```

```

    mam;
    rdm;
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls [diffserv-te](#)]
- [edit logical-systems *logical-system-name* protocols mpls [diffserv-te](#)]



NOTE: If you change the bandwidth model on an ingress router, all the LSPs enabled on the router are taken down and resigaled.

Configuring Traffic Engineering Classes

Configuring traffic engineering classes is optional. [Table 5 on page 317](#) shows the default values for everything in the traffic engineering class matrix. The default mapping is expressed in terms of the default forwarding classes defined in the CoS configuration.

Table 5: Default Values for the Traffic Engineering Class Matrix

Traffic Engineering Class	Class Type	Queue	Priority
te0	ct0	0	7
te1	ct1	1	7
te2	ct2	2	7
te3	ct3	3	7
te4	ct0	0	0
te5	ct1	1	0
te6	ct2	2	0
te7	ct3	3	0

If you want to override the default mappings, you can configure traffic engineering classes 0 through 7. For each traffic engineering class, you configure a class type (or queue) from 0 through 3. For each class type, you configure a priority from 0 through 7.

To configure traffic engineering classes explicitly, include the **te-class-matrix** statement:

```

te-class-matrix {
  tnumber {
    priority priority;
    traffic-class {
      ctnumber priority priority;
    }
  }
}

```

```

    }
  }

```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls [diffserv-te](#)]
- [edit logical-systems *logical-system-name* protocols mpls [diffserv-te](#)]

The following example shows how to configure traffic engineering class **te0** with class type **ct1** and a priority of 4:

```

[edit protocols mpls diffserv-te]
te-class-matrix {
    te0 traffic-class ct1 priority 4;
}

```



NOTE: If you explicitly configure a value for one of the traffic engineering classes, all the default values in the traffic engineering class matrix are dropped.

When you explicitly configure traffic engineering classes, you must also configure a bandwidth model; otherwise, the configuration commit operation fails.

Requirements and Limitations for the Traffic Engineering Class Matrix

When you configure a traffic engineering class matrix, be aware of the following requirements and limitations:

- A mapping configuration is local and affects only the router on which it is configured. It does not affect other systems participating in the differentiated services domain. However, for a Differentiated Services domain to function properly, you need to configure the same traffic engineering class matrix on all the routers participating in the same domain.
- When explicitly configuring traffic engineering classes, you must configure the classes in sequence (**te0**, **te1**, **te2**, **te3**, and so on); otherwise, the configuration commit operation fails.

The first traffic engineering class you configure must be **te0**; otherwise, the configuration commit operation fails.

Configuring Class of Service for DiffServ-Aware Traffic Engineering

To configure DiffServ-aware traffic engineering, you must also configure class of service. The following example illustrates a class-of-service configuration that would allocate 25 percent of the link bandwidth to each class:

```

class-of-service {
    interfaces {
        all {
            scheduler-map simple-map;
        }
    }
}

```

```
    }  
  }  
  scheduler-maps {  
    simple-map {  
      forwarding-class assured-forwarding scheduler simple_sched;  
      forwarding-class best-effort scheduler simple_sched;  
      forwarding-class network-control scheduler simple_sched;  
      forwarding-class expedited-forwarding scheduler simple_sched;  
    }  
  }  
  schedulers {  
    simple_sched {  
      transmit-rate percent 25;  
      buffer-size percent 25;  
    }  
  }  
}
```

For more information on how to configure class of service, see the *Class of Service Feature Guide for Routing Devices*.

LSP Bandwidth Oversubscription Overview

LSPs are established with bandwidth reservations configured for the maximum amount of traffic you expect to traverse the LSP. Not all LSPs carry the maximum amount of traffic over their links at all times. For example, even if the bandwidth for link A has been completely reserved, actual bandwidth might still be available but not currently in use. This excess bandwidth can be used by allowing other LSPs to also use link A, oversubscribing the link. You can oversubscribe the bandwidth configured for individual class types or specify a single value for all of the class types using an interface.

You can use oversubscription to take advantage of the statistical nature of traffic patterns and to permit higher utilization of links.

The following examples describe how you might use bandwidth oversubscription and undersubscription:

- Use oversubscription on class types where peak periods of traffic do not coincide in time.
- Use oversubscription of class types carrying best-effort traffic. You take the risk of temporarily delaying or dropping traffic in exchange for making better utilization of network resources.
- Give different degrees of oversubscription or undersubscription of traffic for the different class types. For instance, you configure the subscription for classes of traffic as follows:
 - Best effort—**ct0 1000**
 - Voice—**ct3 1**

When you undersubscribe a class type for a multiclass LSP, the total demand of all RSVP sessions is always less than the actual capacity of the class type. You can use undersubscription to limit the utilization of a class type.

The bandwidth oversubscription calculation occurs on the local router only. Because no signaling or other interaction is required from other routers in the network, the feature can be enabled on individual routers without being enabled or available on other routers which might not support this feature. Neighboring routers do not need to know about the oversubscription calculation, they rely on the IGP.

The following sections describe the types of bandwidth oversubscription available in the Junos OS:

- [LSP Size Oversubscription on page 320](#)
- [LSP Link Size Oversubscription on page 320](#)
- [Class Type Oversubscription and Local Oversubscription Multipliers on page 320](#)

LSP Size Oversubscription

For LSP size oversubscription, you simply configure less bandwidth than the peak rate expected for the LSP. You also might need to adjust the configuration for automatic policers. Automatic policers manage the traffic assigned to an LSP, ensuring that it does not exceed the configured bandwidth values. LSP size oversubscription requires that the LSP can exceed its configured bandwidth allocation.

Policing is still possible. However, the policer must be manually configured to account for the maximum bandwidth planned for the LSP, rather than for the configured value.

LSP Link Size Oversubscription

You can increase the maximum reservable bandwidth on the link and use the inflated values for bandwidth accounting. Use the **subscription** statement to oversubscribe the link. The configured value is applied to all class type bandwidth allocations on the link. For more information about link size oversubscription, see [“Configuring the Bandwidth Subscription Percentage for LSPs” on page 323](#).

Class Type Oversubscription and Local Oversubscription Multipliers

Local oversubscription multipliers (LOMs) allow different oversubscription values for different class types. LOMs are useful for networks where the oversubscription ratio needs to be configured differently on different links and where oversubscription values are required for different classes. You might use this feature to oversubscribe class types handling best-effort traffic, but use no oversubscription for class types handling voice traffic. An LOM is calculated locally on the router. No information related to an LOM is signaled to other routers in the network.

An LOM is configurable on each link and for each class type. The per-class type LOM allows you to increase or decrease the oversubscription ratio. The per-class-type LOM is factored into all local bandwidth accounting for admission control and IGP advertisement of unreserved bandwidths.

The LOM calculation is tied to the bandwidth model (MAM, extended MAM, and Russian dolls) used, because the effect of oversubscription across class types must be accounted for accurately.



NOTE: All LOM calculations are performed by the Junos OS and require no user intervention.

The formulas related to the oversubscription of class types are described in the following sections:

- [Class Type Bandwidth and the LOM on page 321](#)
- [LOM Calculation for the MAM and Extended MAM Bandwidth Models on page 321](#)
- [LOM Calculation for the Russian Dolls Bandwidth Model on page 322](#)
- [Example: LOM Calculation on page 322](#)

Class Type Bandwidth and the LOM

The following formula expresses the relationship between the bandwidth of the class type and the LOM. The normalized bandwidth of the class type (N_B) is equal to the reserved bandwidth of the class type (R_B) divided by the LOM of the class type (L_C):

$$N_B = R_B / L_C$$

When calculating available bandwidth, you need to subtract the normalized bandwidth from the relevant bandwidth constraint.



NOTE: When using an LOM, values advertised for the available bandwidth might be larger than the bandwidth constraint values. However, the values advertised in the maximum link bandwidth advertisement are not affected by local oversubscription.

LOM Calculation for the MAM and Extended MAM Bandwidth Models

The following formulas show how the LOM is calculated for the MAM and extended MAM bandwidth models.

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

Or

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

where:

- LOM_c—LOM for class type *c*.
- BC_c—Bandwidth constraint for class type *c*.
- CT_c—Class type *c*.
- TE-Class(*i*) <---> (CT_c , preemption *p*) in the configured TE-Class mapping.

LOM Calculation for the Russian Dolls Bandwidth Model

The following formulas show how the LOM is calculated for the Russian dolls bandwidth model:

$$\begin{aligned} \text{Unreserved TE-Class (i)} &= \text{LOM}_c \times \text{MIN} [\\ &[\text{BC}_c - \text{SUM (Normalized (CT}_b, q)) }] \text{ for } q \leq p \text{ and } c \leq b \leq 7, \\ &\dots \\ &[\text{BC}_0 - \text{SUM (Normalized (CT}_b, q)) }] \text{ for } q \leq p \text{ and } 0 \leq b \leq 7, \\ &] \end{aligned}$$

where:

- LOM_c—LOM for class type *c*.
- BC_c—Bandwidth constraint for class type *c*.
- TE-Class(*i*) <--->(CT_c , preemption *p*) in the configured TE-Class mapping.

Note that the impact of an LSP on the unreserved bandwidth of a class type does not depend only on the LOM for that class type—it also depends on the LOM for the class type of the LSP.

Example: LOM Calculation

The following example illustrates how an LOM calculation is made for four classes of traffic: **ct0**, **ct1**, **ct2**, and **ct3**.

The class types have been assigned the following values:

ct0 = 40
ct1 = 30
ct2 = 20
ct3 = 10

These class type values yield the following bandwidth constraints:

BC0 = (ct3 + ct2 + ct1 + ct0) = 100
BC1 = (ct3 + ct2 + ct1) = 60
BC2 = (ct3 + ct2) = 30
BC3 = (ct3) = 10

LSPs from class type **ct0** can take up to 100 percent of bandwidth on the link. LSPs from class type **ct1** can take up to 60 percent of the bandwidth on the link, and so on.

If you assume for this example that the class types have the following LOM values:

LOM(ct0) = 8
LOM(ct1) = 4
LOM(ct2) = 2
LOM(ct3) = 1

In the absence of any other reservation, LSPs from class type **ct0** can take up to 800 percent of the available bandwidth ($8 \times 100 = 800$). In the absence of any other reservation, LSPs from class type **ct1** can take up to 240 percent of the available bandwidth ($4 \times 60 = 240$). and so on.

The maximum amount of bandwidth that can be reserved is:

ct0 = LOM(ct0) x BC0 = 800
ct1 = LOM(ct1) x BC1 = 240
ct2 = LOM(ct2) x BC2 = 60
ct3 = LOM(ct3) x BC3 = 10

For the undersubscribed class type **ct3**, the maximum reservable bandwidth is the same as the bandwidth constraint. For the overbooked class types, these values are not the values of the bandwidth constraint-taking into account the oversubscription for each class type separately. The oversubscription per class type in the sum is not taken into account because ultimately the entire bandwidth constraint can be filled with the bandwidth reservation of just one class type, so you have to account for that class type's bandwidth oversubscription only.

When calculating the available bandwidth for **CTc**, you need to express reservations from other classes as if they were from **CTc**. The reservation from class **ctx** is normalized with the LOM of **ctx**, but it is then multiplied by the LOM of **CTc**.

For the previous example, assume that **LSP1** has class type **ct3** configured with bandwidth of 10 and a priority of 0.

The values for the reservable bandwidth will be:

ct0 = $8 \times (100 - 10) = 720$
ct1 = $4 \times \min((100-10), (60-10)) = 200$
ct2 = $2 \times \min((100-10), (60-10), (30-10)) = 40$
ct3 = $1 \times \min((100-10), (60-10), (30-10), (10-10)) = 0$

These numbers can be rationalized as follows: the normalized reservation is 10 percent. If this bandwidth came from class type **ct0**, it would be equivalent to an overbooked reservation of 80 percent. You can see that 720 percent ($800 - 80 = 720$) of the bandwidth remains available for other LSPs.

Configuring the Bandwidth Subscription Percentage for LSPs

By default, RSVP allows all of a class type's bandwidth (100 percent) to be used for RSVP reservations. When you oversubscribe a class type for a multiclass LSP, the aggregate demand of all RSVP sessions is allowed to exceed the actual capacity of the class type.

If you want to oversubscribe or undersubscribe all of the class types on an interface using the same percentage bandwidth, configure the percentage using the **subscription** statement:

subscription percentage;

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

To undersubscribe or oversubscribe the bandwidth for each class type, configure a percentage for each class type (**ct0**, **ct1**, **ct2**, and **ct3**) option for the **subscription** statement. When you oversubscribe a class type, an LOM is applied to calculate the actual bandwidth reserved. See “[Class Type Oversubscription and Local Oversubscription Multipliers](#)” on [page 320](#) for more information.

```
subscription {
  ct0 percentage;
  ct1 percentage;
  ct2 percentage;
  ct3 percentage;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

percentage is the percentage of class type bandwidth that RSVP allows to be used for reservations. It can be a value from 0 through 65,000 percent. If you specify a value greater than 100, you are oversubscribing the interface or class type.

The value you configure when you oversubscribe a class type is a percentage of the class type bandwidth that can actually be used. The default subscription value is 100 percent.

You can use the **subscription** statement to disable new RSVP sessions for one or more class types. If you configure a percentage of 0, no new sessions (including those with zero bandwidth requirements) are permitted for the class type.

Existing RSVP sessions are not affected by changing the subscription factor. To clear an existing session, issue the **clear rsvp session** command. For more information on the **clear rsvp session** command, see the [CLI Explorer](#).

Constraints on Configuring Bandwidth Subscription

Be aware of the following issues when configuring bandwidth subscription:

- If you configure bandwidth constraints at the **[edit class-of-service interface *interface-name*]** hierarchy level, they override any bandwidth configuration you specify at the **[edit protocols rsvp interface *interface-name* bandwidth]** hierarchy level for Diffserv-TE. Also note that either of the CoS or RSVP bandwidth constraints can override the interface hardware bandwidth constraints.
- If you configure a bandwidth subscription value for a specific interface that differs from the value configured for all interfaces (by including different values for the **subscription** statement at the **[edit protocols rsvp interface *interface-name*]** and **[edit protocols rsvp interface all]** hierarchy levels), the interface-specific value is used for that interface.
- You can configure subscription for each class type only if you also configure a bandwidth model. If no bandwidth model is configured, the commit operation fails with the following error message:

```

user@host# commit check
[edit protocols rsvp interface all]
'subscription'
RSVP: Must have a diffserv-te bandwidth model configured when configuring
subscription per traffic class.
error: configuration check-out failed

```

- You cannot include the **subscription** statement both in the configuration for a specific class type and the configuration for the entire interface. The commit operation fails with the following error message:

```

user@host# commit check
[edit protocols rsvp interface all]
'subscription'
  RSVP: Cannot configure both link subscription and per traffic class
subscription.
error: configuration check-out failed

```

Configuring LSPs for DiffServ-Aware Traffic Engineering

You must configure the Differentiated Services domain (see [“Configuring Routers for DiffServ-Aware Traffic Engineering” on page 315](#)) before you can enable DiffServ-aware traffic engineering for LSPs. The Differentiated Services domain provides the underlying class types and corresponding traffic engineering classes that you reference in the LSP configuration. The traffic engineering classes must be configured consistently on each router participating in the Differentiated Services domain for the LSP to function properly.



NOTE: You must configure either MAM or RDM as the bandwidth model when you configure DiffServ-aware traffic engineering for LSPs. See [“Configuring the Bandwidth Model” on page 316](#).

The actual data transmitted over this Differentiated Services domain is carried by an LSP. Each LSP relies on the EXP bits of the MPLS packets to enable DiffServ-aware traffic engineering. Each LSP can carry traffic for a single class type.

All the routers participating in the LSP must be Juniper Networks routers running Junos OS Release 6.3 or later. The network can include routers from other vendors and Juniper Networks routers running earlier versions of the Junos OS. However, the DiffServ-aware traffic engineering LSP cannot traverse these routers.



NOTE: You cannot simultaneously configure multiclass LSPs and DiffServ-aware traffic engineering LSPs on the same router.

To enable DiffServ-aware traffic engineering for LSPs, you need to configure the following:

- [Configuring Class of Service for the Interfaces on page 326](#)
- [Configuring IGP on page 326](#)
- [Configuring Traffic-Engineered LSPs on page 326](#)

- [Configuring Policing for LSPs on page 327](#)
- [Configuring Fast Reroute for Traffic-Engineered LSPs on page 327](#)

Configuring Class of Service for the Interfaces

The existing class-of-service (CoS) infrastructure ensures that traffic that is consistently marked receives the scheduling guarantees for its class. The classification, marking, and scheduling necessary to accomplish this are configured using the existing Junos OS CoS features.



NOTE: The Junos OS does not support CoS on ATM interfaces.

For information about how to configure CoS, see the *Class of Service Feature Guide for Routing Devices*.

Configuring IGP

You can configure either IS-IS or OSPF as the IGP. The IS-IS and OSPF configurations for routers supporting LSPs are standard. For information about how to configure these protocols, see the *Junos OS Routing Protocols Library for Routing Devices*.

Configuring Traffic-Engineered LSPs

You configure an LSP by using the standard LSP configuration statements and procedures. To configure DiffServ-aware traffic engineering for the LSP, specify a class type bandwidth constraint by including the **bandwidth** statement:

```
label-switched-path lsp-name {  
  bandwidth {  
    ctnumber bps;  
  }  
}
```

For a list of hierarchy levels at which you can include the **bandwidth** statement, see the statement summary sections for this statement.

If you do not specify a bandwidth for a class type, **ct0** is automatically specified as the queue for the LSP. You can configure only one class type for each LSP, unlike multiclass LSPs.

The class type statements specify bandwidth (in bits per second) for the following classes:

- **ct0**—Bandwidth reserved for class 0
- **ct1**—Bandwidth reserved for class 1
- **ct2**—Bandwidth reserved for class 2
- **ct3**—Bandwidth reserved for class 3

You can configure setup and holding priorities for an LSP, but the following restrictions apply:

- The combination of class and priority must be one of the configured traffic engineering classes. The default setup priority is 7 and the default holding priority is 0.
- Configuring an invalid combination of class type and priority causes the commit operation to fail.
- Automatic bandwidth allocation is not supported. If you configure automatic bandwidth allocation, the commit operation fails.
- LSPs configured with the **bandwidth** statement but without specifying a class type use the default class type **ct0**.
- For migration issues, see Internet draft draft-ietf-tewg-diff-te-07.txt.

Configuring Policing for LSPs

Policing allows you to control the amount of traffic forwarded through a particular LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. You can configure multiple policers for each LSP.

For information about how to configure a policer for an LSP, see [“Configuring Policers for LSPs” on page 347](#).

Configuring Fast Reroute for Traffic-Engineered LSPs

You can configure fast reroute for traffic engineered LSPs (LSPs carrying a single class of traffic). It is also possible to reserve bandwidth on the detour path for the class of traffic when fast reroute is enabled. The same class type number is used for both the traffic engineered LSP and its detour.

If you configure the router to reserve bandwidth for the detour path, a check is made to ensure that the link is capable of handling DiffServ-aware traffic engineering and for CoS capability before accepting it as a potential detour path. Unsupported links are not used.

You can configure the amount of bandwidth to reserve for detours using either the **bandwidth** statement or the **bandwidth-percent** statement. You can only configure one of these statements at a time. If you do not configure either the **bandwidth** statement or the **bandwidth-percent** statement, the default setting is to not reserve bandwidth for the detour path (the bandwidth guarantee will be lost if traffic is switched to the detour).

When you configure the **bandwidth** statement, you can specify the specific amount of bandwidth (in bits per second [bps]) you want to reserve for the detour path. For information, see [“Configuring Fast Reroute” on page 226](#).

The **bandwidth-percent** statement allows you to specify the bandwidth of the detour path as a percentage of the bandwidth configured for the protected path. For example, if you configure 100 millions bps of bandwidth for the protected path and configure 20 for the **bandwidth-percent** statement, the detour path will have 20 million bps of bandwidth reserved for its use.

To configure the percent of bandwidth used by the detour path based on the bandwidth of the protected path, include the **bandwidth-percent** statement:

```
bandwidth-percent percentage;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name* **fast-reroute**]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name* **fast-reroute**]

Configuring Multiclass LSPs

A multiclass LSP is an LSP configured to reserve bandwidth for multiple class types and also carries the traffic for these class types. The differentiated service behavior is determined by the EXP bits.

You must configure the Differentiated Services domain (see “[Configuring Routers for DiffServ-Aware Traffic Engineering](#)” on page 315) before you can enable a multiclass LSP. The Differentiated Services domain provides the underlying class types and corresponding traffic engineering classes that you reference in a multiclass LSP configuration. The traffic engineering classes must be configured consistently on each router participating in the Differentiated Services domain for the multiclass LSP to function properly.



NOTE: You must configure extended MAM as the bandwidth model when you configure multiclass LSPs. See “[Configuring the Bandwidth Model](#)” on page 316.

All the routers participating in a multiclass LSP must be Juniper Networks routers running Junos OS Release 6.2 or later. The network can include routers from other vendors and Juniper Networks routers running earlier versions of the Junos OS. However, the multiclass LSP cannot traverse these routers.

To enable multiclass LSPs, you need to configure the following:

- [Configuring Class of Service for the Interfaces](#) on page 328
- [Configuring the IGP](#) on page 329
- [Configuring Class-Type Bandwidth Constraints for Multiclass LSPs](#) on page 329
- [Configuring Policing for Multiclass LSPs](#) on page 330
- [Configuring Fast Reroute for Multiclass LSPs](#) on page 330

Configuring Class of Service for the Interfaces

The existing class-of-service infrastructure ensures that traffic that is consistently marked receives the scheduling guarantees for its class. The classification, marking, and scheduling necessary to consistently mark traffic are configured with the existing Junos OS CoS features.



NOTE: The Junos OS does not support ATM interfaces.

For information about how to configure CoS, see the *Class of Service Feature Guide for Routing Devices*.

Configuring the IGP

You can configure either IS-IS or OSPF. The IS-IS and OSPF configurations for routers supporting multiclass LSPs are standard. For information about how to configure these protocols, see the *Junos OS Routing Protocols Library for Routing Devices*.

Configuring Class-Type Bandwidth Constraints for Multiclass LSPs

You configure a multiclass LSP by using the standard LSP configuration statements and procedures. To configure an LSP as a multiclass LSP, specify the class type bandwidth constraints by including the **bandwidth** statement:

```
bandwidth {
  ct0 bps;
  ct1 bps;
  ct2 bps;
  ct3 bps;
}
```

For a list of hierarchy levels at which you can include the **bandwidth** statement, see the statement summary sections for these statements.

The class type statements specify bandwidth (in bits per second) for the following classes:

- **ct0**—Bandwidth reserved for class 0
- **ct1**—Bandwidth reserved for class 1
- **ct2**—Bandwidth reserved for class 2
- **ct3**—Bandwidth reserved for class 3

For example, to configure 50 megabytes of bandwidth for class type 1 and 30 megabytes of bandwidth for class type 2, include the **bandwidth** statement as follows:

```
[edit protocols mpls]
label-switched-path traffic-class {
  bandwidth {
    ct1 50M;
    ct2 30M;
  }
}
```

You cannot configure a bandwidth for a class type and also configure a bandwidth at the **[edit protocols mpls label-switched-path *lsp-name* bandwidth]** hierarchy level. For example, the following configuration cannot be committed:

```
[edit protocols mpls]
label-switched-path traffic-class {
```

```
bandwidth {  
    20M;  
    ct1 10M;  
}  
}
```

You can configure setup and holding priorities for a multiclass LSP, but the following restrictions apply:

- The setup and holding priorities apply to all classes for which bandwidth is requested.
- The combination of class and priority must be one of the configured traffic engineering classes. The default traffic engineering class configuration results in multiclass LSPs that cannot preempt and cannot be preempted. The default setup priority is 7 and the default holding priority is 0.
- Configuring an invalid combination of class type and priority causes the commit operation to fail.
- Automatic bandwidth allocation is not supported for multiclass LSPs. If you configure automatic bandwidth allocation, the commit operation fails.
- LSPs configured with the **bandwidth** statement but without specifying a class type use the default class type **ct0**.

Configuring Policing for Multiclass LSPs

Policing allows you to control the amount of traffic forwarded through a particular multiclass LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. You can configure multiple policers for each multiclass LSP. You can also enable automatic policing for multiclass LSPs.

For information about how to configure a policer for a multiclass LSP, see [“Configuring Policers for LSPs” on page 347](#) and [“Configuring Automatic Policers” on page 349](#).

Configuring Fast Reroute for Multiclass LSPs

You can enable fast reroute for multiclass LSPs. The bandwidth guarantees for the class types can be carried over to the detour path in case the primary path of the multiclass LSP fails. The same traffic class types configured for the primary multiclass LSP are also signaled for the detour LSP.

The bandwidth guarantee for the detour path is a percentage of the bandwidth configured for the class types of the primary path. For example, you configure a value of 50 percent for the detour path and the protected LSP carries traffic for class types CT0 through CT3. The detour path is signaled with the same class types (CT0 through CT3) but with 50 percent of the bandwidth configured for the protected LSP.

If you configure the router to reserve bandwidth for the detour path, a check is made to ensure that the link is capable of handling DiffServ-aware traffic engineering, that all of the traffic class types needed are available, and for CoS capability before accepting it as a potential detour path. Unsupported links are not used.

The bandwidth percentage for fast reroute is signaled from the ingress router to the egress router. All of the intermediate devices must complete their own CSPF computations and signaling.

When you configure the **bandwidth-percent** statement, the detour path bandwidth is computed by multiplying by the bandwidth configured for the primary multiclass LSP. For information about how to configure the bandwidth for the multiclass LSP, see [“Configuring Traffic-Engineered LSPs” on page 326](#).

To configure the percentage of bandwidth used by the detour path based on the bandwidth of the protected path, include the **bandwidth-percent** statement:

```
bandwidth-percent percentage;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name fast-reroute*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name fast-reroute*]

CHAPTER 8

Configuring Miscellaneous MPLS Properties

- [Configuring the Maximum Number of MPLS Labels on page 334](#)
- [Configuring MPLS to Pop the Label on the Ultimate-Hop Router on page 335](#)
- [Advertising Explicit Null Labels to BGP Peers on page 336](#)
- [Configuring Traffic Engineering for LSPs on page 336](#)
- [Enabling Interarea Traffic Engineering on page 339](#)
- [Enabling Inter-AS Traffic Engineering for LSPs on page 340](#)
- [Configuring MPLS to Gather Statistics on page 342](#)
- [Configuring System Log Messages and SNMP Traps for LSPs on page 344](#)
- [Configuring MPLS Firewall Filters and Policers on page 345](#)
- [Configuring MPLS Rewrite Rules on page 353](#)
- [Configuring BFD for MPLS IPv4 LSPs on page 354](#)
- [BFD-Triggered Local Repair for Rapid Convergence on page 357](#)
- [Pinging LSPs on page 359](#)
- [Tracing MPLS and LSP Packets and Operations on page 361](#)
- [Configuring Link State Distribution Using BGP on page 362](#)
- [Example: Configuring Link State Distribution Using BGP on page 364](#)
- [Dynamic Bandwidth Management Using Container LSP Overview on page 382](#)
- [Configuring Dynamic Bandwidth Management Using Container LSP on page 409](#)
- [Example: Configuring Dynamic Bandwidth Management Using Container LSP on page 413](#)
- [Configuring On-Demand Loss and Delay Measurement on page 438](#)
- [Example: Configuring On-Demand Loss and Delay Measurement on page 439](#)
- [Configuring Pro-Active Loss and Delay Measurements on page 448](#)
- [Example: Configuring Pro-active Loss and Delay Measurements for Bidirectional MPLS LSPs on page 450](#)

Configuring the Maximum Number of MPLS Labels

For interfaces that you configure for MPLS applications, you can set the maximum number of labels upon which MPLS can operate.

By default, the maximum number of labels is three. You can change the maximum to four labels or five labels for applications that require four or five labels. For example, suppose you configure a two-tier carrier-of-carriers VPN service for customers who provide VPN service. A carrier-of-carrier VPN is a two-tiered relationship between a provider carrier (Tier 1 ISP) and a customer carrier (Tier 2 ISP). In a carrier-of-carrier VPN, the provider carrier provides a VPN backbone network for the customer carrier. The customer carrier in turn provides Layer 3 VPN service to its end customers. The customer carrier sends labeled traffic to the provider carrier to deliver it to the next hop on the other side of the provider carrier's network. This scenario requires a three-label stack: one label for the provider carrier VPN, another label for the customer carrier VPN, and a third label for the transport route.

If you add fast reroute service, the PE routers in the provider carrier's network must be configured to support a fourth label (the reroute label). If the customer carrier is using LDP as its signaling protocol and the provider carrier is using RSVP, the provider carrier must support LDP over RSVP tunnel service. This additional service requires an additional label, for a total of five labels.

To the customer carrier, the router it uses to connect to the provider carrier's VPN is a PE router. However, the provider carrier views this device as a CE router.

[Table 6 on page 334](#) summarizes the label requirements.

Table 6: Sample Scenarios for Using 3, 4, or 5 MPLS Labels

Number of Labels Required	Scenarios
3	Carrier-of-carriers VPN or a VPN with two labels and fast reroute
4	Combination of carrier-of-carriers and fast reroute
5	Carrier-of-carriers with fast reroute and the customer carrier running LDP, with the provider carrier running RSVP

The system reserves label space when you configure the maximum number of labels on the interface. When you configure features that require MPLS labels, the label push is automatic. You do not need to explicitly push the labels. The transport route can be a static, LDP-signaled, or RSVP-signaled LSP.

This feature is supported on the following routers:

- MX Series 3D Universal Edge Router
- M120 Multiservice Edge Router
- M320 Multiservice Edge Router with Enhanced III FPCs

- M7i Multiservice Edge Router and M10i Multiservice Edge Router with Enhanced Compact Forwarding Engine Board (CFEB-E)
- T640, T1600, TX Matrix, and TX Matrix Plus routers with Enhanced Scaling FPC1, Enhanced Scaling FP2, Enhanced Scaling FPC3, and Enhanced Scaling FPC4.

To configure and monitor the maximum number of labels:

1. Specify the maximum on the logical interface. Apply this configuration to the carrier's PE routers.

```
[edit interfaces ge-0/1/3 unit 0 family mpls]
user@switch# set maximum-labels 5
```

2. Verify the configuration.

```
[edit system]
user@switch# show interfaces ge-0/1/3.0
Logical interface ge-0/1/3.0 (Index 77) (SNMP ifIndex 507)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Protocol mpls, MTU: 1480, Maximum labels: 5
  Flags: Is-Primary
```

The command output includes the **Maximum labels: 5** field under the logical interface unit 0.

Related Documentation

- [Fast Reroute Overview on page 46](#)
- [Tunneling LDP LSPs in RSVP LSPs Overview on page 523](#)
- *Junos VPNs Configuration Guide* for a carrier-of-carriers configuration example

Configuring MPLS to Pop the Label on the Ultimate-Hop Router

You can control the label value advertised on the egress router of a label-switched path (LSP). The default advertised label is label 3 (Implicit Null Label). If label 3 is advertised, the penultimate-hop router removes the label and sends the packet to the egress router. By enabling ultimate-hop popping, label 0 (IPv4 Explicit Null Label) is advertised. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label.



NOTE: Juniper Networks routers queue packets based on the incoming label. Routers from other vendors might queue packets differently. Keep this in mind when working with networks containing routers from multiple vendors.

To configure MPLS to pop the label on the ultimate-hop router, include the **explicit-null** statement:

```
explicit-null;
```

You can configure this statement at the following hierarchy levels:

- [\[edit protocols mpls\]](#)
- [\[edit logical-systems *logical-system-name* protocols mpls\]](#)

**Related
Documentation**

- [MPLS Label Overview on page 24](#)
- [MPLS Label Allocation on page 26](#)

Advertising Explicit Null Labels to BGP Peers

For the IPv4 (**inet**) family only, BGP peers in a routing group can send an explicit NULL label for a set of connected routes (direct and loopback routes) for the inet labeled-unicast and inet6 labeled-unicast NLRI. By default, peers advertise label 3 (implicit NULL). If the **explicit-null** statement is enabled, peers advertise label 0 (explicit NULL). The explicit NULL labels ensures that labels are always present on packets traversing an MPLS network. If the implicit NULL label is used, the penultimate hop router removes the label and sends the packet as a plain IP packet to the egress router. This might cause issues in queuing the packet properly on the penultimate hop router if the penultimate hop is another vendor's router. Some other vendors queue packets based on the CoS bits in the outgoing label rather than the incoming label.

To advertise an explicit null label, include the following statements in the configuration:

```
family inet {
  labeled-unicast {
    aggregate-label {
      community community-name;
    }
    explicit-null {
      connected-only;
    }
  }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The **connected-only** statement is required to advertise explicit null labels.

To verify that the explicit NULL label is being advertised for connected routes, use the **show route advertising-protocol bgp *neighbor-address*** command.

**Related
Documentation**

- [Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router on page 629](#)
- [Configuring RSVP to Pop the Label on the Ultimate-Hop Router on page 489](#)

Configuring Traffic Engineering for LSPs

When you configure an LSP, a host route (a 32-bit mask) is installed in the ingress router toward the egress router; the address of the host route is the destination address of the

LSP. The **bgp** option for the **traffic engineering** statement at the **[edit protocols mpls]** hierarchy level is enabled by default (you can also explicitly configure the **bgp** option), allowing only BGP to use LSPs in its route calculations. The other **traffic-engineering** statement options allow you to alter this behavior in the master routing instance. This functionality is not available for specific routing instances. Also, you can enable only one of the **traffic-engineering** statement options (**bgp**, **bgp-igp**, **bgp-igp-both-ribs**, or **mpls-forwarding**) at a time.



NOTE: Enabling or disabling any of the **traffic-engineering** statement options causes all the MPLS routes to be removed and then reinserted into the routing tables.

You can configure OSPF and traffic engineering to advertise the LSP metric in summary link-state advertisements (LSAs) as described in the section “[Advertising the LSP Metric in Summary LSAs](#)” on page 339.

The following sections describe how to configure traffic engineering for LSPs:

- [Using LSPs for Both BGP and IGP Traffic Forwarding on page 337](#)
- [Using LSPs for Forwarding in Virtual Private Networks on page 337](#)
- [Using RSVP and LDP Routes for Forwarding but Not Route Selection on page 338](#)
- [Advertising the LSP Metric in Summary LSAs on page 339](#)

Using LSPs for Both BGP and IGP Traffic Forwarding

You can configure BGP and the IGPs to use LSPs for forwarding traffic destined for egress routers by including the **bgp-igp** option for the **traffic-engineering** statement. The **bgp-igp** option causes all inet.3 routes to be moved to the inet.0 routing table.

On the ingress router, include **bgp-igp** option for the **traffic-engineering** statement:

```
traffic-engineering bgp-igp;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**



NOTE: The **bgp-igp** option for the **traffic-engineering** statement cannot be configured for VPN). VPNs require that routes be in the inet.3 routing table.

Using LSPs for Forwarding in Virtual Private Networks

VPNs require that routes remain in the inet.3 routing table to function properly. For VPNs, configure the **bgp-igp-both-ribs** option of the **traffic-engineering** statement to cause BGP and the IGPs to use LSPs for forwarding traffic destined for egress routers. The

bgp-igp-both-ribs option installs the ingress routes in both the inet.0 routing table (for IPv4 unicast routes) and the inet.3 routing table (for MPLS path information).

On the ingress router, include the **traffic-engineering bgp-igp-both-ribs** statement:

```
traffic-engineering bgp-igp-both-ribs;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

When you use the **bgp-igp-both-ribs** statement, the routes from the inet.3 table get copied into the inet.0 table. The copied routes are LDP-sigaled or RSVP-sigaled, and are likely to have a lower preference than other routes in inet.0. Routes with a lower preference are more likely to be chosen as the active routes. This can be a problem because routing policies only act upon active routes. To prevent this problem, use the **mpls-forwarding** option instead.

Using RSVP and LDP Routes for Forwarding but Not Route Selection

If you configure the **bgp-igp** or **bgp-igp-both-ribs** options for the **traffic-engineering** statement, high-priority LSPs can supersede IGP routes in the inet.0 routing table. IGP routes might no longer be redistributed since they are no longer the active routes.

If you configure the **mpls-forwarding** option for the **traffic-engineering** statement, LSPs are used for forwarding but are excluded from route selection. These routes are added to both the inet.0 and inet.3 routing tables. LSPs in the inet.0 routing table are given a low preference when the active route is selected. However, LSPs in the inet.3 routing table are given a normal preference and are therefore used for selecting forwarding next hops.

When you activate the **mpls-forwarding** option, routes whose state is **ForwardingOnly** are preferred for forwarding even if their preference is lower than that of the currently active route. To examine the state of a route, execute a **show route detail** command.

To use LSPs for forwarding but exclude them from route selection, include the **mpls-forwarding** option for the **traffic-engineering** statement:

```
traffic-engineering mpls-forwarding;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

When you configure the **mpls-forwarding** option, IGP shortcut routes are copied to the inet.0 routing table only.

Unlike the **bgp-igp-both-ribs** option, the **mpls-forwarding** option allows you to use the LDP-sigaled and RSVP-sigaled routes for forwarding, and keep the BGP and IGP routes active for routing purposes so that routing policies can act upon them.

For example, suppose a router is running BGP and it has a BGP route of 10.10.10.1/32 that it needs to send to another BGP speaker. If you use the **bgp-igp-both-ribs** option, and your router also has a label-switched-path (LSP) to 10.10.10.1, the MPLS route for 10.10.10.1 becomes active in the inet.0 routing table. This prevents your router from advertising the 10.10.10.1 route to the other BGP router. On the other hand, if you use the **mpls-forwarding** option instead of the **bgp-igp-both-ribs** option, the 10.10.10.1/32 BGP route is advertised to the other BGP speaker, and the LSP is still used to forward traffic to the 10.10.10.1 destination.

Advertising the LSP Metric in Summary LSAs

You can configure MPLS and OSPF to treat an LSP as a link. This configuration allows other routers in the network to use this LSP. To accomplish this goal, you need to configure MPLS and OSPF traffic engineering to advertise the LSP metric in summary LSAs.

For MPLS, include the **traffic-engineering bgp-igp** and **label-switched-path** statements:

```
traffic-engineering bgp-igp;
label-switched-path lsp-name {
    to address;
}
```

You can include these statements at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

For OSPF, include the **lsp-metric-into-summary** statement:

```
lsp-metric-into-summary;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols ospf traffic-engineering shortcuts]**
- **[edit logical-systems *logical-system-name* protocols ospf traffic-engineering shortcuts]**

For more information about OSPF traffic engineering, see the *Junos OS Routing Protocols Library for Routing Devices*.

Enabling Interarea Traffic Engineering

The Junos OS can signal a contiguous traffic-engineered LSP across multiple OSPF areas. The LSP signaling must be done using either nesting or contiguous signaling, as described in RFC 4206, *Label-Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)*. However, contiguous signaling support is limited to just basic signaling. Reoptimization is not supported with contiguous signaling.

The following describes some of the interarea traffic engineering features:

- Interarea traffic engineering can be enabled when the loose-hop area border routers (ABRs) are configured on the ingress router using CSPF for the Explicit Route Object (ERO) calculation within an OSPF area. ERO expansion is completed on the ABRs.

- Interarea traffic engineering can be enabled when CSPF is enabled, but without ABRs specified in the LSP configuration on the ingress router (ABRs can be automatically designated).
- Differentiated Services (DiffServ) traffic engineering is supported as long as the class type mappings are uniform across multiple areas.

To enable interarea traffic engineering, include the **expand-loose-hop** statement in the configuration for each LSP transit router:

expand-loose-hop;

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

Enabling Inter-AS Traffic Engineering for LSPs

Generally, traffic engineering is possible for LSPs that meet the following conditions:

- Both ends of the LSP are in the same OSPF area or at the same IS-IS level.
- The two ends of the LSP are in different OSPF areas within the same autonomous system (AS). LSPs that end in different IS-IS levels are not supported.
- The two ends of an explicit-path LSP are in different OSPF ASs and the autonomous system border routers (ASBRs) are configured statically as the loose hops supported on the explicit-path LSP. For more information, see [“Configuring Explicit-Path LSPs” on page 278](#).

Without statically defined ASBRs on LSPs, traffic engineering is not possible between one routing domain, or AS, and another. However, when the ASs are under the control of single service provider, it is possible in some cases to have traffic engineered LSPs span the ASs and dynamically discover the OSPF ASBRs linking them (IS-IS is not supported with this feature).

Inter-AS traffic engineered LSPs are possible as long as certain network requirements are met, none of the limiting conditions apply, and OSPF passive mode is configured with EBGP. Details are provided in the following sections:

- [Inter-AS Traffic Engineering Requirements on page 340](#)
- [Inter-AS Traffic Engineering Limitations on page 341](#)
- [Configuring OSPF Passive TE Mode on page 342](#)

Inter-AS Traffic Engineering Requirements

The proper establishment and functioning of inter-AS traffic engineered LSPs depend on the following network requirements, all of which must be met:

- All ASs are under control of a single service provider.
- OSPF is used as the routing protocol within each AS, and EBGp is used as the routing protocol between the ASs.
- ASBR information is available inside each AS.
- EBGp routing information is distributed by OSPF, and an IBGP full mesh is in place within each AS.
- Transit LSPs are *not* configured on the inter-AS links, but *are* configured between entry and exit point ASBRs on each AS.
- The EBGp link between ASBRs in different ASs is a direct link and must be configured as a passive traffic engineering link under OSPF. The remote link address itself, not the loopback or any other link address, is used as the remote node identifier for this passive link. For more information about OSPF passive traffic engineering mode configuration, see [“Configuring OSPF Passive TE Mode” on page 342](#).

In addition, the address used for the remote node of the OSPF passive traffic engineering link must be the same as the address used for the EBGp link. For more information about OSPF and BGP in general, see the *Junos OS Routing Protocols Library for Routing Devices*.

Inter-AS Traffic Engineering Limitations

Only LSP hierarchical, or nested, signaling is supported for inter-AS traffic engineered LSPs. Only point-to-point LSPs are supported (there is no point-to-multipoint support).

In addition, the following limitations apply. Any one of these conditions is sufficient to render inter-AS traffic engineered LSPs impossible, even if the above requirements are met.

- The use of multihop BGP is not supported.
- The use of policers or topologies that prevent BGP routes from being known inside the AS is not supported.
- Multiple ASBRs on a LAN between EBGp peers are not supported. Only one ASBR on a LAN between EBGp peers is supported (others ASBRs can exist on the LAN, but cannot be advertised).
- Route reflectors or policies that hide ASBR information or prevent ASBR information from being distributed inside the ASs are not supported.
- Bidirectional LSPs are not supported (LSPs are unidirectional from the traffic engineering perspective).
- Topologies with both inter-AS and intra-AS paths to the same destination are not supported.

In addition, several features that are routine with all LSPs are not supported with inter-AS traffic engineering:

- Admin group link colors are not supported.
- Secondary standby is not supported.

- Reoptimization is not supported.
- Crankback on transit routers is not supported.
- Diverse path calculation is not supported.
- Graceful restart is not supported.

These lists of limitations or unsupported features with inter-AS traffic engineered LSPs are not exhaustive.

Configuring OSPF Passive TE Mode

Ordinarily, interior routing protocols such as OSPF are not run on links between ASs. However, for inter-AS traffic engineering to function properly, information about the inter-AS link, in particular, the address on the remote interface, must be made available inside the AS. This information is not normally included either in EBGP reachability messages or in OSPF routing advertisements.

To flood this link address information within the AS and make it available for traffic engineering calculations, you must configure OSPF passive mode for traffic engineering on each inter-AS interface. You must also supply the remote address for OSPF to distribute and include in the traffic engineering database.

To configure OSPF passive mode for traffic engineering on an inter-AS interface, include the **passive** statement for the link at the **[edit protocols ospf area *area-id* interface *interface-name*]** hierarchy level:

```
passive {
  traffic-engineering {
    remote-node-id ip-address; /* IP address at far end of inter-AS link */
  }
}
```

OSPF must be properly configured on the router. The following example configures the inter-AS link **so-1/1/0** to distribute traffic engineering information with OSPF within the AS. The remote IP address is **192.168.207.2**.

```
[edit protocols ospf area 0.0.0.0]
interface so-1/1/0 {
  unit 0 {
    passive {
      traffic-engineering {
        remote-node-id 192.168.207.2;
      }
    }
  }
}
```

Configuring MPLS to Gather Statistics

You can configure MPLS so that it periodically gathers traffic statistics about all MPLS sessions, including transit sessions, by configuring the **statistics** statement. You must configure the **statistics** statement if you want to collect MPLS traffic statistics using SNMP polling of MPLS Management Information Bases (MIBs).

To enable or disable MPLS statistics collection, include the **statistics** statement:

```
statistics {
  auto-bandwidth;
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval seconds;
  no-transit-statistics;
  transit-statistics-polling;
}
```

You can configure these statements at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

The default interval is 300 seconds.

If you configure the **file** option, the statistics are placed in a file, with one entry per LSP. During the specified interval, the following information is recorded in this file:

- The number of packets, number of bytes, packets per second, and bytes per second transmitted by each LSP. Feature parity for the display of packet and byte statistics for sub-LSPs of a point-to-multipoint LSP on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.
- The percent of bandwidth transmitted over a given LSP in relation to the bandwidth percentage configured for that LSP. If no bandwidth is configured for an LSP, 0 percent is recorded in the percentage column.

At the end of each periodic report, a summary shows the current time, total number of sessions, number of sessions read, number of sessions ignored, and read errors, if any. Ignored sessions are typically those not in the up state or those with a reserved (0 through 15) incoming label (typically the egress point of an LSP). The reason for a read error appears on the same line as the entry for the LSP on which the error occurred. Gathering statistics is an unreliable process; occasional read errors might affect their accuracy. Sample output follows:

lsp6	0 pkt	0 Byte	0 pps	0 Bps	0
lsp5	0 pkt	0 Byte	0 pps	0 Bps	0
lsp6.1	34845 pkt	2926980 Byte	1049 pps	88179 Bps	132
lsp5.1	0 pkt	0 Byte	0 pps	0 Bps	0
lsp4	0 pkt	0 Byte	0 pps	0 Bps	0
Dec 7 17:28:38 Total 6 sessions: 5 success, 0 fail, 1 ignored					

Related Documentation

- [Configuring Automatic Bandwidth Allocation for LSPs on page 257](#)

Configuring System Log Messages and SNMP Traps for LSPs

Whenever an LSP makes a transition from up to down, or down to up, and whenever an LSP switches from one active path to another, the ingress router generates a system log message and sends an SNMP trap. The following shows a sample system log message:

```
MPLS lsp sheep1 up on primary(any) Route 192.168.1.1 192.168.1.2 192.168.1.3
MPLS lsp sheep1 change on primary(any) Route 192.168.1.1 192.168.1.2 192.168.1.3
MPLS lsp sheep1 down on primary(any)
MPLS lsp sheep1 up on secondary(any) Route 192.168.1.1 192.168.1.2 192.168.1.3
MPLS lsp sheep1 change on secondary(any) to primary(any), Route 192.168.1.1
192.168.1.2 192.168.1.3
```

For information about the MPLS SNMP traps and the proprietary MPLS MIBs, see the *Network Management Administration Guide for Routing Devices*.

To generate system log messages for LSPs, include the **syslog** option to the **log-updown** statement:

```
log-updown {
    syslog;
}
```

To generate SNMP traps for LSPs, include the **trap** option to the **log-updown** statement:

```
log-updown {
    trap;
}
```

To generate SNMP traps whenever an LSP path goes down, include the **trap-path-down** option to the **log-updown** statement:

```
log-updown {
    trap-path-down;
}
```

To generate SNMP traps whenever an LSP path comes up, include the **trap-path-up** option to the **log-updown** statement:

```
log-updown {
    trap-path-up;
}
```

To disable the generation of system log messages, include the **no-syslog** option to the **log-updown** statement:

```
log-updown {
    no-syslog;
}
```

To disable the generation of SNMP traps, include the **no-trap** statement:

```
no-trap {
    mpls-lsp-traps;
    rfc3812-traps;
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls log-updown]**
- **[edit logical-systems *logical-system-name* protocols mpls log-updown]**

For scalability reasons, only the ingress router generates SNMP traps. By default, MPLS issues traps for all configured LSPs. If you have many LSPs, the number of traps can become quite large. To disable the generation of SNMP traps, configure the **no-trap** statement.

The **no-trap** statement also includes the following options which allow you to block certain categories of MPLS SNMP traps:

- **mpls-lsp-traps**—Blocks the MPLS LSP traps defined in the **jnx-mpls.mib**, but allows the **rfc3812.mib** traps.
- **rfc-3812-traps**—Blocks the traps defined in the **rfc3812.mib**, but allows the MPLS LSP traps defined in the **jnx-mpls.mib**.

Configuring MPLS Firewall Filters and Policers

You can configure an MPLS firewall filter to count packets based on the EXP bits for the top-level MPLS label in a packet. You can also configure policers for MPLS LSPs.

The following sections discuss MPLS firewall filters and policers:

- [Configuring MPLS Firewall Filters on page 345](#)
- [Examples: Configuring MPLS Firewall Filters on page 346](#)
- [Configuring Policers for LSPs on page 347](#)
- [Example: Configuring an LSP Policer on page 348](#)
- [Configuring Automatic Policers on page 349](#)
- [Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets on page 352](#)

Configuring MPLS Firewall Filters

You can configure an MPLS firewall filter to count packets based on the EXP bits for the top-level MPLS label in a packet. You can then apply this filter to a specific interface. You can also configure a policer for the MPLS filter to police (that is, rate-limit) the traffic on the interface to which the filter is attached. You cannot apply MPLS firewall filters to Ethernet (fxp0) or loopback (lo0) interfaces.

You can configure an MPLS firewall filter on the M Series Multiservice Edge Routers and the T Series Core Routers.

You can configure the following match criteria attributes for MPLS filters at the **[edit firewall family mpls filter *filter-name* term *term-name* from]** hierarchy level:

- **exp**
- **exp-except**

These attributes can accept EXP bits in the range 0 through 7. You can configure the following choices:

- A single EXP bit—for example, **exp 3**;
- Several EXP bits—for example, **exp 0, 4**;
- A range of EXP bits—for example, **exp [0-5]**;

If you do not specify a match criterion (that is, you do not configure the **from** statement and use only the **then** statement with the **count** action keyword), all the MPLS packets passing through the interface on which the filter is applied will be counted.

You also can configure any of the following action keywords at the **[edit firewall family mpls filter *filter-name* term *term-name* then]** hierarchy level:

- **count**
- **accept**
- **discard**
- **next**
- **policer**

For more information about how to configure firewall filters, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*. For more information about how to configure interfaces, see the *Junos OS Network Interfaces Library for Routing Devices* and the *Junos OS Services Interfaces Library for Routing Devices*.

Examples: Configuring MPLS Firewall Filters

The following examples illustrate how you might configure an MPLS firewall filter and then apply the filter to an interface. This filter is configured to count MPLS packets with EXP bits set to either 0 or 4.

The following shows a configuration for an MPLS firewall filter:

```
[edit firewall]
family mpls {
  filter expf {
    term expt0 {
      from {
        exp 0,4;
      }
      then {
        count counter0;
        accept;
      }
    }
  }
}
```

The following shows how to apply the MPLS firewall filter to an interface:

```
[edit interfaces]
```



```

so-0/0/0 {
  mtu 4474;
  encapsulation ppp;
  sonet-options {
    fcs 32;
  }
  unit 0 {
    point-to-point;
    family mpls {
      filter {
        input expf;
        output expf;
      }
    }
  }
}

```

The MPLS firewall filter is applied to the input and output of an interface (see the **input** and **output** statements in the preceding example).

Configuring Policers for LSPs

MPLS LSP policing allows you to control the amount of traffic forwarded through a particular LSP. Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. LSP policing is supported on regular LSPs, LSPs configured with DiffServ-aware traffic engineering, and multiclass LSPs. You can configure multiple policers for each multiclass LSP. For regular LSPs, each LSP policer is applied to all of the traffic traversing the LSP. The policer's bandwidth limitations become effective as soon as the total sum of traffic traversing the LSP exceeds the configured limit.

You configure the multiclass LSP and DiffServ-aware traffic engineering LSP policers in a filter. The filter can be configured to distinguish between the different class types and apply the relevant policer to each class type. The policers distinguish between class types based on the EXP bits.

You configure LSP policers under the **family any** filter. The **family any** filter is used because the policer is applied to traffic entering the LSP. This traffic might be from different families: IPv6, MPLS, and so on. You do not need to know what sort of traffic is entering the LSP, as long as the match conditions apply to all types of traffic.

You can configure only those match conditions that apply across all types of traffic. The following are the supported match conditions for LSP policers:

- **forwarding-class**
- **packet-length**
- **interface**
- **interface-set**

To enable a policer on an LSP, first you need to configure a policing filter and then include it in the LSP configuration. For information about how to configure policers, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

To configure a policer for an LSP, specify a filter by including the **filter** option to the **policing** statement:

```
policing {  
    filter filter-name;  
}
```

You can include the **policing** statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit protocols mpls **static-label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **static-label-switched-path** *lsp-name*]

LSP Policer Limitations

When configuring MPLS LSP policers, be aware of the following limitations:

- LSP policers are supported for packet LSPs only.
- LSP policers are supported for unicast next hops only. Multicast next hops are not supported.
- LSP policers are not supported on aggregated interfaces.
- The LSP policer runs before any output filters.
- Traffic sourced from the Routing Engine (for example, ping traffic) does not take the same forwarding path as transit traffic. This type of traffic cannot be policed.
- LSP policers work on all T Series routers and on M Series routers that have the Internet Processor II application-specific integrated circuit (ASIC).



NOTE: Starting with Junos OS Release 12.2R2, on T Series routers only, you can configure an LSP policer for a specific LSP to be shared across different protocol family types. To do so, you must configure the *logical-interface-policer* statement at the [edit firewall policer *policer-name*] hierarchy level.

Example: Configuring an LSP Policer

The following example shows how you can configure a policing filter for an LSP:

```
[edit firewall]  
policer police-ct1 {  
    if-exceeding {  
        bandwidth-limit 50m;  
        burst-size-limit 1500;  
    }  
    then {  
        discard;  
    }  
}
```

```

    }
  }
  policer police-ct0 {
    if-exceeding {
      bandwidth-limit 200m;
      burst-size-limit 1500;
    }
    then {
      discard;
    }
  }
  family any {
    filter bar {
      term discard-ct0 {
        then {
          policer police-ct0;
          accept;
        }
      }
    }
    term discard-ct1 {
      then {
        policer police-ct1;
        accept;
      }
    }
  }
}

```

Configuring Automatic Policers

Automatic policing of LSPs allows you to provide strict service guarantees for network traffic. Such guarantees are especially useful in the context of Differentiated Services for traffic engineered LSPs, providing better emulation for ATM wires over an MPLS network. For more information about Differentiated Services for LSPs, see [“DiffServ-Aware Traffic Engineering Introduction” on page 310](#).

Differentiated Services for traffic engineered LSPs allow you to provide differential treatment to MPLS traffic based on the EXP bits. To ensure these traffic guarantees, it is insufficient to simply mark the traffic appropriately. If traffic follows a congested path, the requirements might not be met.

LSPs are guaranteed to be established along paths where enough resources are available to meet the requirements. However, even if the LSPs are established along such paths and are marked properly, these requirements cannot be guaranteed unless you ensure that no more traffic is sent to an LSP than there is bandwidth available.

It is possible to police LSP traffic by manually configuring an appropriate filter and applying it to the LSP in the configuration. However, for large deployments it is cumbersome to configure thousands of different filters. Configuration groups cannot solve this problem either, since different LSPs might have different bandwidth requirements, requiring different filters. To police traffic for numerous LSPs, it is best to configure automatic policers.

When you configure automatic policers for LSPs, a policer is applied to all of the LSPs configured on the router. However, you can disable automatic policing on specific LSPs.



NOTE: When you configure automatic policers for DiffServ-aware traffic engineering LSP, GRES is not supported.



NOTE: You cannot configure automatic policing for LSPs carrying CCC traffic.

The following sections describe how to configure automatic policers for LSPs:

- [Configuring Automatic Policers for LSPs on page 350](#)
- [Configuring Automatic Policers for DiffServ-Aware Traffic Engineering LSPs on page 351](#)
- [Configuring Automatic Policers for Point-to-Multipoint LSPs on page 351](#)
- [Disabling Automatic Policing on an LSP on page 352](#)
- [Example: Configuring Automatic Policing for an LSP on page 352](#)

Configuring Automatic Policers for LSPs

To configure automatic policers for standard LSPs (neither DiffServ-aware traffic engineered LSPs nor multiclass LSPs), include the **auto-policing** statement with either the **class all** *policer-action* option or the **class ct0** *policer-action* option:

```
auto-policing {
    class all policer-action;
    class ct0 policer-action;
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

You can configure the following policer actions for automatic policers:

- **drop**—Drop all packets.
- **loss-priority-high**—Set the packet loss priority (PLP) to high.
- **loss-priority-low**—Set the PLP to low.

These policer actions are applicable to all types of LSPs. The default policer action is to do nothing.

Automatic policers for LSPs police traffic based on the amount of bandwidth configured for the LSPs. You configure the bandwidth for an LSP using the **bandwidth** statement at the **[edit protocols mpls label-switched-path *lsp-path-name*]** hierarchy level. If you have enabled automatic policers on a router, change the bandwidth configured for an LSP, and commit the revised configuration, the change does not take effect on the active

LSPs. To force the LSPs to use the new bandwidth allocation, issue a **clear mpls lsp** command.



NOTE: You cannot configure automatic policers for LSPs that traverse aggregated interfaces or Multilink Point-to-Point Protocol (MLPPP) interfaces.

Configuring Automatic Policers for DiffServ-Aware Traffic Engineering LSPs

To configure automatic policers for DiffServ-aware traffic engineering LSPs and for multiclass LSPs, include the **auto-policing** statement:

```
auto-policing {
  class all policer-action;
  class ctnumber policer-action;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

You include either the **class all *policer-action*** statement or a **class *ctnumber* *policer-action*** statement for each of one or more classes (you can configure a different policer action for each class). For a list of the actions that you can substitute for the ***policer-action*** variable, see “[Configuring Automatic Policers for LSPs](#)” on page 350. The default policer action is to do nothing.



NOTE: You cannot configure automatic policers for LSPs that traverse aggregated interfaces or MLPPP interfaces.

Configuring Automatic Policers for Point-to-Multipoint LSPs

You can configure automatic policers for point-to-multipoint LSPs by including the **auto-policing** statement with either the **class all *policer-action*** option or the **class *ct0* *policer-action*** option. You only need to configure the **auto-policing** statement on the primary point-to-multipoint LSP (for more information on primary point-to-multipoint LSPs, see “[Configuring the Primary Point-to-Multipoint LSP](#)” on page 283). No additional configuration is required on the subLSPs for the point-to-multipoint LSP.

Point-to-multipoint automatic policing is applied to all branches of the point-to-multipoint LSP. In addition, automatic policing is applied to any local VRF interfaces that have the same forwarding entry as a point-to-multipoint branch. Feature parity for automatic policers for MPLS point-to-multipoint LSPs on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.

The automatic policer configuration for point-to-multipoint LSPs is identical to the automatic policer configuration for standard LSPs. For more information, see “[Configuring Automatic Policers for LSPs](#)” on page 350.

Disabling Automatic Policing on an LSP

When you enable automatic policing, all of the LSPs on the router or logical system are affected. To disable automatic policing on a specific LSP on a router where you have enabled automatic policing, include the **policing** statement with the **no-auto-policing** option:

```
policing no-auto-policing;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

Example: Configuring Automatic Policing for an LSP

Configure automatic policing for a multiclass LSP, specifying different actions for class types **ct0**, **ct1**, **ct2**, and **ct3**.

```
[edit protocols mpls]
diffserv-te {
  bandwidth-model extended-mam;
}
auto-policing {
  class ct1 loss-priority-low;
  class ct0 loss-priority-high;
  class ct2 drop;
  class ct3 loss-priority-low;
}
traffic-engineering bgp-igp;
label-switched-path sample-lsp {
  to 3.3.3.3;
  bandwidth {
    ct0 11;
    ct1 1;
    ct2 1;
    ct3 1;
  }
}
interface fxp0.0 {
  disable;
}
interface t1-0/5/3.0;
interface t1-0/5/4.0;
```

Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets

You can selectively set the DiffServ code point (DSCP) field of MPLS-tagged IPv4 and IPv6 packets to 0 without affecting output queue assignment, and continue to set the MPLS EXP field according to the configured rewrite table, which is based on forwarding classes. You can accomplish this by configuring a firewall filter for the MPLS-tagged packets.

For instructions on how to write different DSCP and EXP values in MPLS-tagged IP packets, see the *Class of Service Feature Guide for Routing Devices*. For instructions on how to configure firewall filters, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Related Documentation

- [Firewall Filter Match Conditions for MPLS Traffic](#)

Configuring MPLS Rewrite Rules

You can apply a number of different rewrite rules to MPLS packets.

For more information about how to configure statements at the **[edit class-of-service]** hierarchy level, see the *Class of Service Feature Guide for Routing Devices*.

The following sections describe how you can apply rewrite rules to MPLS packets:

- [Rewriting the EXP Bits of All Three Labels of an Outgoing Packet on page 353](#)
- [Rewriting MPLS and IPv4 Packet Headers on page 353](#)

Rewriting the EXP Bits of All Three Labels of an Outgoing Packet

In interprovider, carrier-of-carrier, and complex traffic engineering scenarios, it is sometimes necessary to push three labels on the next hop.

By default, on M Series routers except the M320, the top MPLS EXP label of an outgoing packet is not rewritten when you configure swap-push-push and triple-push operations. You can rewrite the EXP bits of all three labels of an outgoing packet, thereby maintaining the class of service (CoS) of an incoming MPLS or non-MPLS packet.

To push three labels on incoming MPLS packets, include the **exp-swap-push-push default** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
exp-swap-push-push default;
```

To push three labels on incoming non-MPLS packets, include the **exp-push-push-push default** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules]
exp-push-push-push default;
```

For more information about how to configure statements at the **[edit class-of-service]** hierarchy level, see the *Class of Service Feature Guide for Routing Devices*.

Rewriting MPLS and IPv4 Packet Headers

You can apply a rewrite rule to MPLS and IPv4 packet headers simultaneously. This allows you to initialize MPLS EXP and IP precedence bits at LSP ingress. You can configure different rewrite rules depending on whether the traffic is VPN or non-VPN.

To rewrite MPLS and IPv4 packet headers, include the **protocol** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules exp *rewrite-rule-name*]** hierarchy level:

```
[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules exp  
  rewrite-rule-name]  
protocol types;
```

Use the **protocol** statement to specify the types of MPLS packets and packet headers to which to apply the rewrite rule. The MPLS packet can be a standard MPLS packet or an MPLS packet with an IPv4 payload. Specify the type of MPLS packet by using the following options:

- **mpls-any**—Applies the rewrite rule to MPLS packets and writes the code point value to MPLS headers.
- **mpls-inet-both**—Applies the rewrite rule to VPN MPLS packets with IPv4 payloads. Writes the code point value to the MPLS and IPv4 headers in T Series (except T4000 routers) and M320 routers. On M Series routers, except the M320, the **mpls-inet-both** option causes all ingress MPLS LSP packets with IPv4 payloads to be initialized with 000 code points for IP precedence and MPLS EXP values.
- **mpls-inet-both-non-vpn**—Applies the rewrite rule to any non-VPN MPLS packets with IPv4 payloads. Writes the code point value to the MPLS and IPv4 headers in T Series and M320 routers. On M Series routers, except the M320, the **mpls-inet-both-non-vpn** option causes all ingress MPLS LSP packets with IPv4 payloads to be initialized with 000 code points for IP precedence and MPLS EXP values.

For a detailed example on how to configure rewrite rules for MPLS and IPv4 packets and for more information about how to configure class of service, see the *Class of Service Feature Guide for Routing Devices*.

Configuring BFD for MPLS IPv4 LSPs

You can configure Bidirectional Forwarding Detection (BFD) protocol on MPLS IPv4 LSPs as outlined in the Internet draft draft-ietf-bfd-mpls-02.txt, *BFD for MPLS LSPs*. BFD is used as a periodic Operation, Administration, and Maintenance (OAM) feature for LSPs to detect LSP data plane faults. You can configure BFD for LSPs that use either LDP or RSVP as the signaling protocol.

You can also use the LSP **ping** commands to detect LSP data plane faults. However, BFD has a couple of benefits: it requires less computer processing than LSP **ping** commands and can quickly detect faults in large numbers of LSPs (LSP **ping** commands must be issued for each LSP individually). On the other hand, BFD cannot be used to verify the control plane against the data plane at the egress LSR, which is possible when an LSP **ping** echo request is associated with a forwarding equivalence class (FEC).

The BFD failure detection timers are adaptive and can be adjusted to be more or less aggressive. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local

BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

Starting from Junos OS Release 13.2R4, 13.3R2, and 14.1, you can configure the **lsp-ping-interval** statement and the **lsp-ping-multiplier** statement at the **[edit protocols mpls oam]** hierarchy level to set the time interval between LSP ping messages and the number of LSP ping responses respectively after which the Bidirectional Forwarding Detection (BFD) session is brought down.

For configuration instructions for LDP-signaled LSPs, see [“Configuring BFD for LDP LSPs” on page 543](#). For configuration instructions for RSVP-signaled LSPs, see the following section.

Configuring BFD for RSVP-Signaled LSPs

BFD for RSVP supports unicast IPv4 LSPs. When BFD is configured for an RSVP LSP on the ingress router, it is enabled on the primary path and on all standby secondary paths for that LSP. The source IP address for outgoing BFD packets from the egress side of an MPLS BFD session is based on the outgoing interface IP address. You can enable BFD for all LSPs on a router or for specific LSPs. If you configure BFD for a specific LSP, whatever values configured globally for BFD are overridden. The BFD sessions originate only at the ingress router and terminate at the egress router.

An error is logged whenever a BFD session for a path fails. The following example shows how BFD for RSVP LSP log messages might appear:

```
RPD_MPLS_PATH_BFD_UP: MPLS BFD session for path path1 up on LSP R0_to_R3
RPD_MPLS_PATH_BFD_DOWN: MPLS BFD session for path path1 down on LSP R0_to_R3
```

You can configure BFD for all of the RSVP LSPs on the router, a specific LSP, or the primary path of a specific LSP. To configure BFD for RSVP LSPs, include the **oam** and **bfd-liveness-detection** statements.

```
oam {
  bfd-liveness-detection {
    failure-action {
      make-before-break teardown-timeout seconds;
      teardown;
    }
    failure-action teardown;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier detection-time-multiplier;
  }
  lsp-ping-interval time-interval;
  lsp-ping-multiplier multiplier;
}
```

You can configure this statement at the following hierarchy levels:

- `[edit protocols mpls]`
- `[edit protocols mpls label-switched-path lsp-name]`
- `[edit protocols mpls label-switched-path lsp-name primary path-name]`

The **bfd-liveness-detection** statement includes the following options:

- **minimum-interval**—Specifies the minimum transmit and receive interval.
- **minimum-receive-interval**—Specifies the minimum receive interval. The range is from 1 through 255,000 milliseconds.
- **minimum-transmit-interval**—Specifies the minimum transmit interval. The range is from 1 through 255,000 milliseconds.
- **lsp-ping-multiplier**—Specifies the detection time multiplier. The range is from 1 through 255.



NOTE: To avoid triggering false negatives, configure a BFD fault detection time that is longer than the fast reroute time.

You can also configure the **lsp-ping-interval** option to adjust the time interval between LSP pings. The LSP ping command for RSVP-signaled LSPs is **ping mpls rsvp**. For more information on the **ping mpls rsvp** command, see the [CLI Explorer](#).

Configuring a Failure Action for the BFD Session on an RSVP LSP

When the BFD session for an RSVP LSP goes down, the LSP is torn down and resignaled. Traffic can be switched to a standby LSP, or you can simply tear down the LSP path. Any actions performed are logged.

When a BFD session for an RSVP LSP path goes down, you can configure the Junos OS to resignal the LSP path or to simply disable the LSP path. A standby LSP path could be configured to handle traffic while the primary LSP path is unavailable. The router can automatically recover from LSP failures that can be detected by BFD. By default, if a BFD session fails, the event is simply logged.

To enable the Junos OS to tear down an RSVP LSP path in the event of a BFD event, include the **failure-action** statement:

```
failure-action {
  make-before-break teardown-timeout seconds;
  teardown;
}
```

For a list of the hierarchy levels at which you can include this statement, see the statement summary section for this statement.

You can configure either the **teardown** or **make-before-break** options:

- **teardown**—Causes the LSP path to be taken down and resignaled immediately.

- **make-before-break**—Causes the Junos OS to attempt to signal a new LSP path before tearing down the old LSP path. You can also configure the **teardown-timeout** option to automatically tear down the LSP after the time period specified if the attempt to resignal the LSP fails within the **teardown-timeout** interval. If you specify a value of 0 for the **teardown-timeout** interval, the LSP is taken down and resignaled immediately (the same behavior as when you configure the **teardown** option).

To configure a failure action for all of the RSVP LSPs, include the **failure-action** statement at the **[edit protocols mpls oam bfd-liveness-detection]** hierarchy level. To configure a failure action for a specific RSVP LSP, include the **failure-action** statement at the **[edit protocols mpls label-switched-path *lsp-name* oam bfd-liveness-detection]** hierarchy level.

To configure a failure action for a specific primary path, include the **failure-action** statement at the **[edit protocols mpls label-switched path *lsp-name* primary *path-name* oam bfd-liveness-detection]** hierarchy level. To configure a failure action for a specific secondary LSP path, include the **failure-action** statement at the **[edit protocols mpls label-switched-path *lsp-name* secondary *path-name* oam bfd-liveness-detection]** hierarchy level.

BFD-Triggered Local Repair for Rapid Convergence

- [Understanding BFD-Triggered Local Protection on page 357](#)
- [Disabling BFD-Triggered Local Repair on page 358](#)

Understanding BFD-Triggered Local Protection

The time it takes for a network to converge following a link or node failure can vary dramatically based on a number of factors, including network size, the protocols used, and network design. However, while each particular convergence event is different, the process of convergence is essentially consistent. The failure is detected, the failure is reported (flooded) in the network, an alternate path is found for traffic, and the forwarding plane is updated to pass traffic on a new path.

This overview discusses how Bidirectional Forwarding Detection (BFD)-triggered local repair contributes to a quicker restoration time for rapid convergence in an MPLS network.

- [Purpose of BFD-Triggered Local Repair on page 357](#)
- [Configuring BFD-Triggered Local Repair on page 358](#)


Purpose of BFD-Triggered Local Repair

In Junos OS, general MPLS traffic protection for RSVP-signaled label-switched path (LSP) failures is provided by several complementary mechanisms. These protection mechanisms include local protection (fast reroute, link protection, and node-link protection) and path protection (primary and secondary paths). Local protection in conjunction with path protection can provide minimum packet loss for an LSP, and control the way the LSP is rerouted after a failure. Traditionally, both types of protection rely on fast detection of connectivity failure at the physical level. However, for transmission media without fast physical level detection, Junos OS supports BFD and MPLS ping for fast failure detection.

With links between routers, when a route goes down, the routing protocol process recalculates the next best path. When MPLS fast reroute (FRR) is enabled, ifl messages are flooded to all Flexible PIC Concentrators (FPCs). The edge FPC enables the bypass MPLS LSP tunnel. Lastly, all routes are repaired and sent through the bypass MPLS LSP tunnel. The amount of time it takes to repair all routes is proportional to the number of routes.

This repair scenario becomes more difficult when a switch lies between two links. See [Figure 34 on page 358](#).

Figure 34: Topology with BFD-Triggered Local Repair



When a link goes down at the remote end, the failure is not detected at the local end until the interior gateway protocol (IGP) goes down. To wait for the routing protocol process to recalculate the next best path takes too much time.

With BFD-triggered local repair enabled, the Packet Forwarding Engine completes the repair first, using the bypass MPLS LSP tunnel (that is preconfigured and installed), then informs the routing protocol process to recalculate a new route. By doing this, when the primary MPLS LSP tunnel goes down, the FPC can intermittently and immediately divert traffic to the FPC with the bypass MPLS LSP tunnel.

Using local repair in this way achieves a faster restoration time of less than 50 ms.

Configuring BFD-Triggered Local Repair

BFD-triggered local repair is not configurable, but is part of the default configuration.

BFD-triggered local repair works within the legacy Junos OS features MPLS-FRR, BFD for IGP, and loop-free alternates (LFAs).

Disabling BFD-Triggered Local Repair

By default, BFD-triggered local repair is enabled for all routing interfaces. If desired, you can disable BFD-triggered local repair at the `[edit routing-options]` hierarchy level.

Disabling BFD-Triggered Local Repair

To explicitly disable BFD-triggered local repair:

1. Include the `no-bfd-triggered-local-repair` statement at the `[edit routing-options]` hierarchy level:

```
user@host# set no-bfd-triggered-local-repair
```

2. (Optional) Verify your configuration settings before committing them by using the **show routing-options** command.

```
user@host# run show routing-options
```

Confirm your configuration by issuing the **show routing-options** command.

```
user@host# show routing-options
...
no-bfd-triggered-local-repair;
}
```



NOTE: When you disable this feature, you must also restart routing by including the **graceful-restart** statement for the IGP. For example, for OSPF, this is accomplished by including the **graceful-restart** statement at the [edit **protocols ospf**] hierarchy level.

Related Documentation

- [Configuring BFD for LDP LSPs on page 543](#)
- [Configuring Link Protection on Interfaces Used by LSPs on page 502](#)
- [Configuring Fast Reroute on page 226](#)
- [Configuring Graceful Restart for Point-to-Multipoint LSPs on page 305](#)
- *graceful-restart (Protocols OSPF)*

Pinging LSPs

The following sections describe how to use the **ping mpls** command to confirm LSP functioning.

- [Pinging MPLS LSPs on page 359](#)
- [Pinging Point-to-Multipoint LSPs on page 360](#)
- [Pinging the Endpoint Address of MPLS LSPs on page 360](#)
- [Pinging CCC LSPs on page 360](#)
- [Pinging Layer 3 VPNs on page 360](#)
- [Support for LSP Ping and Traceroute Commands Based on RFC 4379 on page 360](#)

Pinging MPLS LSPs

You can ping a specific LSP. Echo requests are sent over the LSP as MPLS packets. The payload is a User Datagram Protocol (UDP) packet forwarded to an address in the 127/8 range (127.0.0.1 by default, this address is configurable) and port 3503. The label and interface information for building and sending this information as an MPLS packet is the same as for standard LSP traffic.

When the echo request arrives at the egress node, the receiver checks the contents of the packet and sends a reply containing the correct return value, by using UDP. The router

sending the echo request waits to receive an echo reply after a timeout of 2 seconds (you cannot configure this value).

You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router to be able to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

To ping an MPLS LSP use the **ping mpls <count count> <ldp <fec>> <rsvp <exp forwarding-class> <lsp-name>>** command. To ping a secondary MPLS LSP, use the **ping mpls <count count> <rsvp <lsp-name>> standby path-name** command. For a detailed description of this command, see the [CLI Explorer](#).



NOTE: The **ping mpls** command is not supported within routing instances.

Pinging Point-to-Multipoint LSPs

To ping a point-to-multipoint LSP, use the **ping mpls rsvp lsp-name multipoint** or **ping mpls rsvp egress address** commands. The **ping mpls rsvp lsp-name multipoint** command returns a list of all of the egress router identifiers and the current status of the point-to-multipoint LSP egress routers. The **ping mpls rsvp lsp-name multipoint egress address** command returns the current status of the specified egress router.

Pinging the Endpoint Address of MPLS LSPs

To determine whether an LSP between two provider edge (PE) routers is up and running, you can ping the endpoint address of the LSP. To ping an MPLS LSP endpoint, use the **ping mpls lsp-end-point address** command. This command tells you what type of LSP (RSVP or LDP) terminates at the address specified and whether that LSP is up or down.

For a detailed description of this command, see the [CLI Explorer](#).

Pinging CCC LSPs

You can ping a specific CCC LSP. The CCC LSP ping command is identical to the one used for MPLS LSPs. The command you use is **ping mpls <count count> <rsvp <lsp-name>>**. You can also ping a secondary standby CCC LSP by using the **ping mpls <count count> <rsvp <lsp-name>> standby path-name** command.

For a detailed description of this command, see the [CLI Explorer](#).

Pinging Layer 3 VPNs

You can use a similar command, **ping mpls l3vpn vpn-name prefix prefix <count count>**, to ping a Layer 3 VPN. For more information about this command, see the *Junos OS VPNs Library for Routing Devices* and the [CLI Explorer](#).

Support for LSP Ping and Traceroute Commands Based on RFC 4379

The Junos OS supports LSP **ping** and **traceroute** commands based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*.

LSP **ping** and **traceroute** commands based on RFC 4379 attempt to trace the path taken by an LSP by relying on MPLS TTL expiration. An LSP can take multiple paths from ingress to egress. This occurs in particular with Equal Cost Multipath (ECMP). The LSP **traceroute** command can trace all possible paths to an LSP node.

Tracing MPLS and LSP Packets and Operations

To trace MPLS and LSP packets and operations, include the **traceoptions** statement:

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

You can specify the following MPLS-specific flags in the MPLS **traceoptions** statement:

- **all**—Trace all operations.
- **connection**—Trace all circuit cross-connect (CCC) activity.
- **connection-detail**—Trace detailed CCC activity.
- **cspf**—Trace CSPF computations.
- **cspf-link**—Trace links visited during CSPF computations.
- **cspf-node**—Trace nodes visited during CSPF computations.
- **error**—Trace MPLS error conditions.
- **graceful-restart**—Trace MPLS graceful restart events.
- **lsping**—Trace LSP ping packets and return codes.
- **nsr-synchronization**—Trace nonstop routing (NSR) synchronization events.
- **nsr-synchronization-detail**—Trace NSR synchronization events in detail.
- **state**—Trace all LSP state transitions.
- **static**—Trace static label-switched path.

When you configure trace options to track an MPLS LSP using the **cspf** option, the CSPF log displays information about the MPLS LSP using the term “generalized MPLS” (GMPLS). For example, a message in the CSPF log might state that the “link passes GMPLS constraints”. Generalized MPLS (GMPLS) is a superset of MPLS, so this message is normal and does not affect proper MPLS LSP operation.

For general information about tracing and global tracing options, see the *Junos OS Routing Protocols Library for Routing Devices*.

Configuring Link State Distribution Using BGP

You can enable distribution of topology information across multiple areas and autonomous systems (ASs) by extending the BGP protocol to carry link-state information, which was initially acquired using IGP. The IGP protocols have scaling limitations when it comes to distributing large databases. BGP is not only a more scalable vehicle for carrying multi-area and multi-AS topology information, but also provides the policy controls that can be useful for multi-AS topology distribution. The BGP link-state topology information is used for computing paths for MPLS LSPs spanning multiple domains, such as inter-area TE LSP, and providing a scalable and policy-controlled means for external path computing entities, such as ALTO and PCE, to acquire network topology.

Before you begin:

1. Configure the device interfaces.
2. Configure the router ID and autonomous system number for the device.
3. Configure the following protocols:
 - RSVP
 - MPLS
 - IS-IS
 - OSPF

To enable link-state distribution using BGP:

1. Configure an internal BGP group, and assign the local address and neighbor address for the group.

[edit protocols]

```
user@R1# set bgp group internal-group-name type internal
user@R1# set bgp group internal-group-name local-address ip-address
user@R1# set bgp group internal-group-name neighbor ip-address
```

2. Include the BGP-TE signaling network layer reachability information (NLRI) to the internal BGP group.

[edit protocols]

```
user@R1# set bgp group internal-group-name family traffic-engineering unicast
```

3. Enable export of policy on the device.

[edit protocols]

```
user@R1# set bgp group internal-group-name export second-policy-name
```

4. Configure an external BGP group, and assign the local address and neighbor autonomous system to the group.

[edit protocols]

```
user@R1# set bgp group external-group-name type external
user@R1# set bgp group external-group-name neighbor ip-address local-address
ip-address
user@R1# set bgp group external-group-name neighbor ip-address peer-as as-number
```


5. Include the BGP-TE signaling NLRI to the external BGP group.

```
[edit protocols]
user@R1# set bgp group external-group-name family traffic-engineering unicast
```

6. In configuration mode, go to the following hierarchy level:

```
[edit]
user@R1# edit policy-options
```

7. Configure policies to accept traffic from the BGP-TE NLRI.

```
[edit policy-options]
user@R1# set policy-statement policy-name from family traffic-engineering
user@R1# set policy-statement policy-name then accept
```

```
user@R1# set policy-statement bgp-import-policy term 1 from family traffic-engineering
user@R1# set policy-statement bgp-import-policy term 1 then next-hop self
user@R1# set policy-statement bgp-import-policy term 1 then accept
```

8. On the remote connecting device, configure policy to accept the OSPF and IS-IS traffic.

```
[edit policy-options]
user@R2# set policy-statement bgp-export-policy term 1 from protocol isis
user@R2# set policy-statement bgp-export-policy term 1 from protocol ospf
user@R2# set policy-statement bgp-export-policy term 1 then accept
user@R2# set policy-statement bgp-export-policy term 2 then reject
```

9. Verify and commit the configuration.

For example:

R1

```
[edit protocols]
user@R1# set rsvp interface all
user@R1# set rsvp interface fxp0.0 disable
user@R1# set mpls interface all
user@R1# set mpls interface fxp0.0 disable
user@R1# set bgp group ibgp type internal
user@R1# set bgp group ibgp local-address 10.255.105.141
user@R1# set bgp group ibgp family traffic-engineering unicast
user@R1# set bgp group ibgp export nlri2bgp
user@R1# set bgp group ibgp neighbor 10.255.105.137
user@R1# set bgp group ebgp type external
user@R1# set bgp group ebgp family traffic-engineering unicast
user@R1# set bgp group ebgp neighbor 8.42.1.104 local-address 8.42.1.102
user@R1# set bgp group ebgp neighbor 8.42.1.104 peer-as 65534
user@R1# set isis interface ge-0/0/1.0 passive remote-node-iso 0102.5502.4211
user@R1# set isis interface ge-0/0/1.0 passive remote-node-id 8.42.1.104
user@R1# set ospf traffic-engineering
user@R1# set ospf area 0.0.0.0 interface lo0.0
user@R1# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@R1# set ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering
remote-node-id 8.42.1.104
user@R1# set ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering
remote-node-router-id 10.255.105.139

[edit policy-options]
```

```
user@R1# set policy-statement accept-all from family traffic-engineering
user@R1# set policy-statement accept-all then accept
user@R1# set policy-statement nlri2bgp term 1 from family traffic-engineering
user@R1# set policy-statement nlri2bgp term 1 then next-hop self
user@R1# set policy-statement nlri2bgp term 1 then accept
```

```
[edit]
user@R1# commit
commit complete
```

R2

```
[edit policy-options]
user@R2# set policy-statement accept-all from family traffic-engineering
user@R2# set policy-statement accept-all then accept
user@R2# set policy-statement nlri2bgp term 1 from family traffic-engineering
user@R2# set policy-statement nlri2bgp term 1 then next-hop self
user@R2# set policy-statement nlri2bgp term 1 then accept
user@R2# set policy-statement ted2nlri term 1 from protocol isis
user@R2# set policy-statement ted2nlri term 1 from protocol ospf
user@R2# set policy-statement ted2nlri term 1 then accept
user@R2# set policy-statement ted2nlri term 2 then reject
```

```
[edit]
user@R2# commit
commit complete
```

- Related Documentation**
- [Link-State Distribution Using BGP Overview on page 8](#)
 - [Example: Configuring Link State Distribution Using BGP on page 364](#)

Example: Configuring Link State Distribution Using BGP

This example shows how to configure BGP to carry link-state information across multiple domains, which is used for computing paths for MPLS LSPs spanning multiple domains, such as inter-area TE LSP, and providing a scalable and policy-controlled means for external path computing entities, such as ALTO and PCE, to acquire network topology.

- [Requirements on page 364](#)
- [Overview on page 365](#)
- [Configuration on page 365](#)
- [Verification on page 375](#)

Requirements

This example uses the following hardware and software components:

- Four routers that can be a combination of M Series, MX Series, or T Series routers
- Junos OS Release 14.2 or later running on all the routers

Before you begin:

1. Configure the device interfaces.

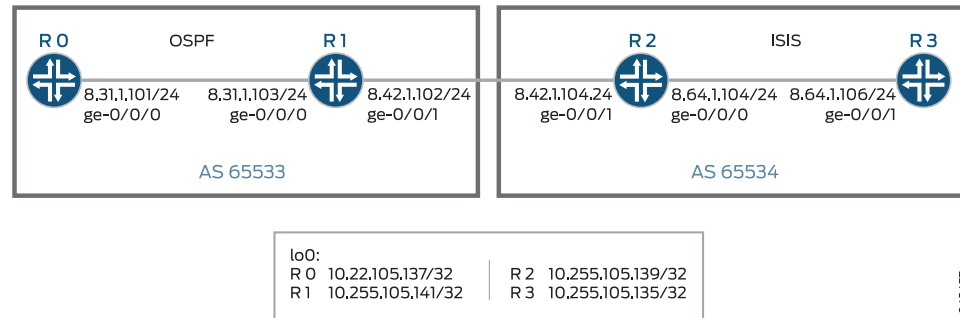
2. Configure the autonomous system numbers and router IDs for the devices.
3. Configure the following protocols:
 - RSVP
 - MPLS
 - BGP
 - IS-IS
 - OSPF

Overview

Starting with Junos OS Release 14.2, a new mechanism to distribute topology information across multiple areas and autonomous systems (ASs) is introduced by extending the BGP protocol to carry link-state information, which was initially acquired using IGP. The IGP protocols have scaling limitations when it comes to distributing large databases. BGP is not only a more scalable vehicle for carrying multi-area and multi-AS topology information, but also provides the policy controls that can be useful for multi-AS topology distribution. The BGP link-state topology information is used for computing paths for MPLS label-switched paths (LSPs) spanning multiple domains, such as inter-area TE LSP, and providing a scalable and policy-controlled means for external path computing entities, such as ALTO and PCE, to acquire network topology.

Topology

Figure 35: Link-State Distribution Using BGP



In Figure 35 on page 365, Routers R0 and R1 and Routers R2 and R3 belong to different autonomous systems. Routers R0 and R1 run OSPF, and Routers R2 and R3 run IS-IS.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
R0    set interfaces ge-0/0/0 unit 0 family inet address 8.31.1.101/24
      set interfaces ge-0/0/0 unit 0 family iso
```

```

set interfaces ge-0/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.105.137/32
set routing-options router-id 10.255.105.137
set routing-options autonomous-system 65533
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls traffic-engineering database export policy accept-all
set protocols mpls cross-credibility-cspf
set protocols mpls label-switched-path to-R3-inter-as to 10.255.105.135
set protocols mpls label-switched-path to-R3-inter-as bandwidth 40m
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.105.137
set protocols bgp group ibgp family traffic-engineering unicast
set protocols bgp group ibgp neighbor 10.255.105.141
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set policy-options policy-statement accept-all from family traffic-engineering
set policy-options policy-statement accept-all then accept

```

```

R1 set interfaces ge-0/0/0 unit 0 family inet address 8.31.1.103/24
set interfaces ge-0/0/0 unit 0 family iso
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 8.42.1.102/24
set interfaces ge-0/0/1 unit 0 family iso
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.105.141/32
set routing-options router-id 10.255.105.141
set routing-options autonomous-system 65533
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 10.255.105.141
set protocols bgp group ibgp family traffic-engineering unicast
set protocols bgp group ibgp export nlri2bgp
set protocols bgp group ibgp neighbor 10.255.105.137
set protocols bgp group ebgp type external
set protocols bgp group ebgp family traffic-engineering unicast
set protocols bgp group ebgp neighbor 8.42.1.104 local-address 8.42.1.102
set protocols bgp group ebgp neighbor 8.42.1.104 peer-as 65534
set protocols isis interface ge-0/0/1.0 passive remote-node-iso 0102.5502.4211
set protocols isis interface ge-0/0/1.0 passive remote-node-id 8.42.1.104
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering
remote-node-id 8.42.1.104
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering
remote-node-router-id 10.255.105.139
set policy-options policy-statement accept-all from family traffic-engineering
set policy-options policy-statement accept-all then accept
set policy-options policy-statement nlri2bgp term 1 from family traffic-engineering

```

```

set policy-options policy-statement nlri2bgp term 1 then next-hop self
set policy-options policy-statement nlri2bgp term 1 then accept

R2  set interfaces ge-0/0/0 unit 0 family inet address 8.64.1.104/24
    set interfaces ge-0/0/0 unit 0 family iso
    set interfaces ge-0/0/0 unit 0 family mpls
    set interfaces ge-0/0/1 unit 0 family inet address 8.42.1.104/24
    set interfaces ge-0/0/1 unit 0 family iso
    set interfaces ge-0/0/1 unit 0 family mpls
    set interfaces lo0 unit 0 family inet address 10.255.105.139/32
    set interfaces lo0 unit 0 family iso
    set routing-options router-id 10.255.105.139
    set routing-options autonomous-system 65534
    set protocols rsvp interface all
    set protocols rsvp interface fxp0.0 disable
    set protocols mpls traffic-engineering database import policy ted2nlri
    set protocols mpls interface all
    set protocols mpls interface fxp0.0 disable
    set protocols bgp group ebgp type external
    set protocols bgp group ebgp family traffic-engineering unicast
    set protocols bgp group ebgp export nlri2bgp
    set protocols bgp group ebgp peer-as 65533
    set protocols bgp group ebgp neighbor 8.42.1.102
    set protocols isis level 1 disable
    set protocols isis interface ge-0/0/0.0
    set protocols isis interface ge-0/0/1.0 passive remote-node-iso 0102.5501.8181
    set protocols isis interface ge-0/0/1.0 passive remote-node-id 8.42.1.102
    set protocols isis interface lo0.0
    set protocols ospf traffic-engineering
    set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering
        remote-node-id 8.42.1.102
    set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering
        remote-node-router-id 10.255.105.141
    set policy-options policy-statement accept-all from family traffic-engineering
    set policy-options policy-statement accept-all then accept
    set policy-options policy-statement nlri2bgp term 1 from family traffic-engineering
    set policy-options policy-statement nlri2bgp term 1 then next-hop self
    set policy-options policy-statement nlri2bgp term 1 then accept
    set policy-options policy-statement ted2nlri term 1 from protocol isis
    set policy-options policy-statement ted2nlri term 1 from protocol ospf
    set policy-options policy-statement ted2nlri term 1 then accept
    set policy-options policy-statement ted2nlri term 2 then reject

R3  set interfaces ge-0/0/0 unit 0 family inet address 8.64.1.106/24
    set interfaces ge-0/0/0 unit 0 family iso
    set interfaces ge-0/0/0 unit 0 family mpls
    set interfaces lo0 unit 0 family inet address 10.255.105.135/32
    set interfaces lo0 unit 0 family iso
    set routing-options router-id 10.255.105.135
    set routing-options autonomous-system 65534
    set protocols rsvp interface all
    set protocols rsvp interface fxp0.0 disable
    set protocols mpls interface all
    set protocols mpls interface fxp0.0 disable
    set protocols bgp group ibgp type internal

```

```

set protocols bgp group ibgp local-address 10.255.105.135
set protocols bgp group ibgp family traffic-engineering unicast
set protocols bgp group ibgp neighbor 10.255.105.139
set protocols isis interface ge-0/0/0.0 level 1 disable
set protocols isis interface lo0.0
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set policy-options policy-statement accept-all from family traffic-engineering
set policy-options policy-statement accept-all then accept
set policy-options policy-statement nlri2bgp term 1 from family traffic-engineering
set policy-options policy-statement nlri2bgp term 1 then next-hop self
set policy-options policy-statement nlri2bgp term 1 then accept
set policy-options policy-statement ted2nlri term 1 from protocol isis
set policy-options policy-statement ted2nlri term 1 from protocol ospf
set policy-options policy-statement ted2nlri term 1 then accept
set policy-options policy-statement ted2nlri term 2 then reject

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure Router R1:

1. Configure the Router R1 interfaces.

```

[edit interfaces]
user@R1# set ge-0/0/0 unit 0 family inet address 8.31.1.103/24
user@R1# set ge-0/0/0 unit 0 family iso
user@R1# set ge-0/0/0 unit 0 family mpls

user@R1# set ge-0/0/1 unit 0 family inet address 8.42.1.102/24
user@R1# set ge-0/0/1 unit 0 family iso
user@R1# set ge-0/0/1 unit 0 family mpls

user@R1# set lo0 unit 0 family inet address 10.255.105.141/32

```

2. Configure the router ID and autonomous system of Router R1.

```

[edit routing-options]
user@R1# set router-id 10.255.105.141
user@R1# set autonomous-system 65533

```

3. Enable RSVP on all the interfaces of Router R1 (excluding the management interface).

```

[edit protocols]
user@R1# set rsvp interface all
user@R1# set rsvp interface fxp0.0 disable

```

4. Enable MPLS on all the interfaces of Router R1 (excluding the management interface).

```

[edit protocols]
user@R1# set mpls interface all
user@R1# set mpls interface fxp0.0 disable

```

5. Configure the BGP group for Router R1 to peer with Router R0, and assign the local address and neighbor address.

```
[edit protocols]
user@R1# set bgp group ibgp type internal
user@R1# set bgp group ibgp local-address 10.255.105.141
user@R1# set bgp group ibgp neighbor 10.255.105.137
```

6. Include the BGP-TE signaling network layer reachability information (NLRI) to the ibgp BGP group.

```
[edit protocols]
user@R1# set bgp group ibgp family traffic-engineering unicast
```

7. Enable export of policy nlri2bgp on Router R1.

```
[edit protocols]
user@R1# set bgp group ibgp export nlri2bgp
```

8. Configure the BGP group for Router R1 to peer with Router R2, and assign the local address and neighbor autonomous system to the ebgp BGP group.

```
[edit protocols]
user@R1# set bgp group ebgp type external
user@R1# set bgp group ebgp neighbor 8.42.1.104 local-address 8.42.1.102
user@R1# set bgp group ebgp neighbor 8.42.1.104 peer-as 65534
```

9. Include the BGP-TE signaling NLRI to the ebgp BGP group.

```
[edit protocols]
user@R1# set bgp group ebgp family traffic-engineering unicast
```

10. Enable passive traffic-engineering on the inter-AS link.

```
[edit protocols]
user@R1# set isis interface ge-0/0/1.0 passive remote-node-iso 0102.5502.4211
user@R1# set isis interface ge-0/0/1.0 passive remote-node-id 8.42.1.104
```

11. Enable OSPF on the interface connecting Router R1 to Router R0 and on the loopback interface of Router R1, and enable traffic engineering capabilities.

```
[edit protocols]
user@R1# set ospf traffic-engineering
user@R1# set ospf area 0.0.0.0 interface lo0.0
user@R1# set ospf area 0.0.0.0 interface ge-0/0/0.0
```

12. Enable passive traffic-engineering on the inter-AS link.

```
[edit protocols]
user@R1# set ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering
remote-node-id 8.42.1.104
user@R1# set ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering
remote-node-router-id 10.255.105.139
```

13. Configure policies to accept traffic from BGP-TE NLRI.

```
[edit policy-options]
user@R1# set policy-statement accept-all from family traffic-engineering
user@R1# set policy-statement accept-all then accept
user@R1# set policy-statement nlri2bgp term 1 from family traffic-engineering
user@R1# set policy-statement nlri2bgp term 1 then next-hop self
user@R1# set policy-statement nlri2bgp term 1 then accept
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 8.31.1.103/24;
    }
    family iso;
    family mpls;
  }
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 8.42.1.102/24;
    }
    family iso;
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.255.105.141/32;
    }
  }
}

user@R1# show routing-options
router-id 10.255.105.141;
autonomous-system 65533;

user@R1# show protocols
rsvp {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
mpls {
  interface all;
  interface fxp0.0 {
    disable;
  }
}
bgp {
  group ibgp {
    type internal;
    local-address 10.255.105.141;
    family traffic-engineering {
      unicast;
    }
  }
}
```



```

        export nlri2bgp;
        neighbor 10.255.105.137;
    }
    group ebgp {
        type external;
        family traffic-engineering {
            unicast;
        }
        neighbor 8.42.1.104 {
            local-address 8.42.1.102;
            peer-as 65534;
        }
    }
}
isis {
    interface ge-0/0/1.0 {
        passive {
            remote-node-iso 0102.5502.4211;
            remote-node-id 8.42.1.104;
        }
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-0/0/0.0;
        interface ge-0/0/1.0 {
            passive {
                traffic-engineering {
                    remote-node-id 8.42.1.104;
                    remote-node-router-id 10.255.105.139;
                }
            }
        }
    }
}
}

user@R1# show policy-options
policy-statement accept-all {
    from family traffic-engineering;
    then accept;
}
policy-statement nlri2bgp {
    term 1 {
        from family traffic-engineering;
        then {
            next-hop self;
            accept;
        }
    }
}

```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure Router R2:

1. Configure the Router R2 interfaces.

```
[edit interfaces]
user@R2# set ge-0/0/0 unit 0 family inet address 8.64.1.104/24
user@R2# set ge-0/0/0 unit 0 family iso
user@R2# set ge-0/0/0 unit 0 family mpls
```

```
user@R2# set ge-0/0/1 unit 0 family inet address 8.42.1.104/24
user@R2# set ge-0/0/1 unit 0 family iso
user@R2# set ge-0/0/1 unit 0 family mpls
```

```
user@R2# set lo0 unit 0 family inet address 10.255.105.139/32
user@R2# set lo0 unit 0 family iso
```

2. Configure the router ID and autonomous system of Router R2.

```
[edit routing-options]
user@R2# set router-id 10.255.105.139
user@R2# set autonomous-system 65534
```

3. Enable RSVP on all the interfaces of Router R2 (excluding the management interface).

```
[edit routing-options]
user@R2# set rsvp interface all
user@R2# set rsvp interface fxp0.0 disable
```

4. Enable MPLS on all the interfaces of Router R2 (excluding the management interface).

```
[edit routing-options]
user@R2# set mpls interface all
user@R2# set mpls interface fxp0.0 disable
```

5. Enable import of traffic engineering database parameters using the ted2nlri policy.

```
[edit protocols]
user@R2# set mpls traffic-engineering database import policy ted2nlri
```

6. Configure the BGP group for Router R2 to peer with Router R1.

```
[edit protocols]
user@R2# set bgp group ebgp type external
```

7. Include the BGP-TE signaling NLRI to the ebgp BGP group.

```
[edit protocols]
user@R2# set bgp group ebgp family traffic-engineering unicast
```

8. Assign the local address and neighbor autonomous system to the ebgp BGP group.

```
[edit protocols]
user@R2# set bgp group ebgp peer-as 65533
user@R2# set bgp group ebgp neighbor 8.42.1.102
```

9. Enable export of policy nlri2bgp on Router R2.

```
[edit protocols]
user@R2# set bgp group ebgp export nlri2bgp
```

10. Enable IS-IS on the interface connecting Router R2 with Router R3 and the loopback interface of Router R2.

```
[edit protocols]
user@R2# set isis level 1 disable
user@R2# set isis interface ge-0/0/0.0
user@R2# set isis interface lo0.0
```

11. Enable only IS-IS advertising on the interface connecting Router R2 with Router R1.

```
[edit protocols]
user@R2# set isis interface ge-0/0/1.0 passive remote-node-iso 0102.5501.8181
user@R2# set isis interface ge-0/0/1.0 passive remote-node-id 8.42.1.102
```

12. Configure traffic engineering capability on Router R2.

```
[edit protocols]
user@R2# set ospf traffic-engineering
```

13. Enable only OSPF advertisements on the interface connecting Router R2 with Router R1.

```
[edit protocols]
user@R2# set ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering
remote-node-id 8.42.1.102
user@R2# set ospf area 0.0.0.0 interface ge-0/0/1.0 passive traffic-engineering
remote-node-router-id 10.255.105.141
```

14. Configure policies to accept traffic from the BGP-TE NLRI.

```
[edit policy-options]
user@R2# set policy-statement accept-all from family traffic-engineering
user@R2# set policy-statement accept-all then accept
user@R2# set policy-statement nlri2bgp term 1 from family traffic-engineering
user@R2# set policy-statement nlri2bgp term 1 then next-hop self
user@R2# set policy-statement nlri2bgp term 1 then accept
user@R2# set policy-statement ted2nlri term 1 from protocol isis
user@R2# set policy-statement ted2nlri term 1 from protocol ospf
user@R2# set policy-statement ted2nlri term 1 then accept
user@R2# set policy-statement ted2nlri term 2 then reject
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 8.64.1.104/24;
    }
    family iso;
    family mpls;
```

```
    }  
  }  
  ge-0/0/1 {  
    unit 0 {  
      family inet {  
        address 8.42.1.104/24;  
      }  
      family iso;  
      family mpls;  
    }  
  }  
  lo0 {  
    unit 0 {  
      family inet {  
        address 10.255.105.139/32;  
      }  
      family iso;  
    }  
  }  
}  
  
user@R2# show routing-options  
router-id 10.255.105.139;  
autonomous-system 65534;  
  
user@R2# show protocols  
rsvp {  
  interface all;  
  interface fxp0.0 {  
    disable;  
  }  
}  
mpls {  
  traffic-engineering {  
    database {  
      import {  
        policy ted2nlri;  
      }  
    }  
  }  
  interface all;  
  interface fxp0.0 {  
    disable;  
  }  
}  
bgp {  
  group ebgp {  
    type external;  
    family traffic-engineering {  
      unicast;  
    }  
    export nlri2bgp;  
    peer-as 65533;  
    neighbor 8.42.1.102;  
  }  
}  
isis {  
  level 1 disable;
```

```

interface ge-0/0/0.0;
interface ge-0/0/1.0 {
  passive {
    remote-node-iso 0102.5501.8181;
    remote-node-id 8.42.1.102;
  }
}
interface lo0.0;
}
ospf {
  traffic-engineering;
  area 0.0.0.0 {
    interface ge-0/0/1.0 {
      passive {
        traffic-engineering {
          remote-node-id 8.42.1.102;
          remote-node-router-id 10.255.105.141;
        }
      }
    }
  }
}
}

user@R2# show policy-options
policy-statement accept-all {
  from family traffic-engineering;
  then accept;
}
policy-statement nlri2bgp {
  term 1 {
    from family traffic-engineering;
    then {
      next-hop self;
      accept;
    }
  }
}
policy-statement ted2nlri {
  term 1 {
    from protocol [ isis ospf ];
    then accept;
  }
  term 2 {
    then reject;
  }
}

```

Verification

Verify that the configuration is working properly.

- [Verifying the BGP Summary Status on page 376](#)
- [Verifying the MPLS LSP Status on page 376](#)

- [Verifying the Ispdist.0 Routing Table Entries on page 377](#)
- [Verifying the Traffic Engineering Database Entries on page 380](#)

Verifying the BGP Summary Status

Purpose Verify that BGP is up and running on Routers R0 and R1.

Action From operational mode, run the **show bgp summary** command.

```
user@R0> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
Ispdist.0
10 10 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.255.105.141 65533 20 14 0 79 5:18
Establ
Ispdist.0: 10/10/10/0
```

From operational mode, run the **show bgp summary** command.

```
user@R1> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
Ispdist.0
10 10 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
8.42.1.104 65534 24 17 0 70 6:43
Establ
Ispdist.0: 10/10/10/0
10.255.105.137 65533 15 23 0 79 6:19
Establ
Ispdist.0: 0/0/0/0
```

Meaning Router R0 is peered with Router R1.

Verifying the MPLS LSP Status

Purpose Verify the status of the MPLS LSP on Router R0.

Action From operational mode, run the **show mpls lsp** command.

```
user@R0> show mpls lsp
Ingress LSP: 1 sessions
  To          From          State Rt P    ActivePath    LSPname
  10.255.105.135 10.255.105.137 Up    0 *          to-R3-inter-as
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning The MPLS LSP from Router R0 to Router R3 is established.

Verifying the Isdist.0 Routing Table Entries

Purpose Verify the Isdist.0 routing table entries on Routers R0, R1, and R2.

Action From operational mode, run the **show route table lsdist.0** command.

```

user@R0> show route table lsdist.0
lsdist.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

NODE { AS:65534 ISO:0102.5502.4211.00 ISIS-L2:0 }/1152
    * [BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
NODE { AS:65534 ISO:0102.5502.4250.00 ISIS-L2:0 }/1152
    * [BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
NODE { AS:65534 ISO:0102.5502.4250.02 ISIS-L2:0 }/1152
    * [BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
NODE { AS:65534 Area:0.0.0.0 IPv4:10.255.105.139 OSPF:0 }/1152
    * [BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
LINK { Local { AS:65534 ISO:0102.5502.4211.00 }. { IPv4:8.42.1.104 } Remote {
AS:65534 ISO:0102.5501.8181.00 }. { IPv4:8.42.1.102 } ISIS-L2:0 }/1152
    * [BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
LINK { Local { AS:65534 ISO:0102.5502.4211.00 }. { IPv4:8.64.1.104 } Remote {
AS:65534 ISO:0102.5502.4250.02 }. { } ISIS-L2:0 }/1152
    * [BGP/170] 00:02:03, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
LINK { Local { AS:65534 ISO:0102.5502.4250.00 }. { IPv4:8.64.1.106 } Remote {
AS:65534 ISO:0102.5502.4250.02 }. { } ISIS-L2:0 }/1152
    * [BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
LINK { Local { AS:65534 ISO:0102.5502.4250.02 }. { } Remote { AS:65534
ISO:0102.5502.4211.00 }. { } ISIS-L2:0 }/1152
    * [BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
LINK { Local { AS:65534 ISO:0102.5502.4250.02 }. { } Remote { AS:65534
ISO:0102.5502.4250.00 }. { } ISIS-L2:0 }/1152
    * [BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0
LINK { Local { AS:65534 Area:0.0.0.0 IPv4:10.255.105.139 }. { IPv4:8.42.1.104 }
Remote { AS:65534 Area:0.0.0.0 IPv4:10.255.105.141 }. { IPv4:8.42.1.102 } OSPF:0
}/1152
    * [BGP/170] 00:17:32, localpref 100, from 10.255.105.141
      AS path: 65534 I, validation-state: unverified
      > to 8.31.1.103 via ge-0/0/0.0

```

From operational mode, run the **show route table lsdist.0** command.

```

user@R1> show route table lsdist.0
lsdist.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```



```

NODE { AS:65534 ISO:0102.5502.4211.00 ISIS-L2:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
    > to 8.42.1.104 via ge-0/0/1.0
NODE { AS:65534 ISO:0102.5502.4250.00 ISIS-L2:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
    > to 8.42.1.104 via ge-0/0/1.0
NODE { AS:65534 ISO:0102.5502.4250.02 ISIS-L2:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
    > to 8.42.1.104 via ge-0/0/1.0
NODE { AS:65534 Area:0.0.0.0 IPv4:10.255.105.139 OSPF:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
    > to 8.42.1.104 via ge-0/0/1.0
LINK { Local { AS:65534 ISO:0102.5502.4211.00 }.{ IPv4:8.42.1.104 } Remote {
AS:65534 ISO:0102.5501.8181.00 }.{ IPv4:8.42.1.102 } ISIS-L2:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
    > to 8.42.1.104 via ge-0/0/1.0
LINK { Local { AS:65534 ISO:0102.5502.4211.00 }.{ IPv4:8.64.1.104 } Remote {
AS:65534 ISO:0102.5502.4250.02 }.{ } ISIS-L2:0 }/1152
    *[BGP/170] 00:02:19, localpref 100
    AS path: 65534 I, validation-state: unverified
    > to 8.42.1.104 via ge-0/0/1.0
LINK { Local { AS:65534 ISO:0102.5502.4250.00 }.{ IPv4:8.64.1.106 } Remote {
AS:65534 ISO:0102.5502.4250.02 }.{ } ISIS-L2:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
    > to 8.42.1.104 via ge-0/0/1.0
LINK { Local { AS:65534 ISO:0102.5502.4250.02 }.{ } Remote { AS:65534
ISO:0102.5502.4211.00 }.{ } ISIS-L2:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
    > to 8.42.1.104 via ge-0/0/1.0
LINK { Local { AS:65534 ISO:0102.5502.4250.02 }.{ } Remote { AS:65534
ISO:0102.5502.4250.00 }.{ } ISIS-L2:0 }/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
    > to 8.42.1.104 via ge-0/0/1.0
LINK { Local { AS:65534 Area:0.0.0.0 IPv4:10.255.105.139 }.{ IPv4:8.42.1.104 }
Remote { AS:65534 Area:0.0.0.0 IPv4:10.255.105.141 }.{ IPv4:8.42.1.102 } OSPF:0
}/1152
    *[BGP/170] 00:18:00, localpref 100
    AS path: 65534 I, validation-state: unverified
    > to 8.42.1.104 via ge-0/0/1.0

```

From operational mode, run the **show route table lsdist.0** command.

```

user@R2> show route table lsdist.0
lsdist.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

NODE { AS:65534 ISO:0102.5502.4211.00 ISIS-L2:0 }/1152
    *[IS-IS/18] 1d 00:24:39
    Fictitious
NODE { AS:65534 ISO:0102.5502.4250.00 ISIS-L2:0 }/1152
    *[IS-IS/18] 00:20:45

```

```

Fictitious
NODE { AS:65534 ISO:0102.5502.4250.02 ISIS-L2:0 }/1152
    *[IS-IS/18] 00:20:45
Fictitious
NODE { AS:65534 Area:0.0.0.0 IPv4:10.255.105.139 OSPF:0 }/1152
    *[OSPF/10] 1d 00:24:39
Fictitious
LINK { Local { AS:65534 ISO:0102.5502.4211.00 }.{ IPv4:8.42.1.104 } Remote {
AS:65534 ISO:0102.5501.8181.00 }.{ IPv4:8.42.1.102 } ISIS-L2:0 }/1152
    *[IS-IS/18] 00:20:58
Fictitious
LINK { Local { AS:65534 ISO:0102.5502.4211.00 }.{ IPv4:8.64.1.104 } Remote {
AS:65534 ISO:0102.5502.4250.02 }.{ } ISIS-L2:0 }/1152
    *[IS-IS/18] 00:02:34
Fictitious
LINK { Local { AS:65534 ISO:0102.5502.4250.00 }.{ IPv4:8.64.1.106 } Remote {
AS:65534 ISO:0102.5502.4250.02 }.{ } ISIS-L2:0 }/1152
    *[IS-IS/18] 00:20:45
Fictitious
LINK { Local { AS:65534 ISO:0102.5502.4250.02 }.{ } Remote { AS:65534
ISO:0102.5502.4211.00 }.{ } ISIS-L2:0 }/1152
    *[IS-IS/18] 00:20:45
Fictitious
LINK { Local { AS:65534 ISO:0102.5502.4250.02 }.{ } Remote { AS:65534
ISO:0102.5502.4250.00 }.{ } ISIS-L2:0 }/1152
    *[IS-IS/18] 00:20:45
Fictitious
LINK { Local { AS:65534 Area:0.0.0.0 IPv4:10.255.105.139 }.{ IPv4:8.42.1.104 }
Remote { AS:65534 Area:0.0.0.0 IPv4:10.255.105.141 }.{ IPv4:8.42.1.102 } OSPF:0
}/1152
    *[OSPF/10] 00:20:57
Fictitious

```

Meaning The routes are appearing in the lsdist.0 routing table.

Verifying the Traffic Engineering Database Entries

Purpose Verify the traffic engineering database entries on Router R0.

Action From operational mode, run the **show ted database** command.

```

user@R0> show ted database
TED database: 5 ISIS nodes 5 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
0102.5501.8168.00(10.255.105.137) Rtr  1046   1     1 OSPF(0.0.0.0)
  To: 8.31.1.101-1, Local: 8.31.1.101, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
0102.5501.8181.00                ---  1033   1     0
0102.5502.4211.00(10.255.105.139) Rtr  3519   2     3 Exported ISIS-L2(1)
  To: 0102.5502.4250.02, Local: 8.64.1.104, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  To: 0102.5501.8181.00, Local: 8.42.1.104, Remote: 8.42.1.102
  Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
                                Exported OSPF(2)
  To: 10.255.105.141, Local: 8.42.1.104, Remote: 8.42.1.102
  Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
0102.5502.4250.00(10.255.105.135) Rtr  1033   1     1 Exported ISIS-L2(1)
  To: 0102.5502.4250.02, Local: 8.64.1.106, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
0102.5502.4250.02                Net   1033   2     2 Exported ISIS-L2(1)
  To: 0102.5502.4211.00(10.255.105.139), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  To: 0102.5502.4250.00(10.255.105.135), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
8.31.1.101-1                     Net   1046   2     2 OSPF(0.0.0.0)
  To: 0102.5501.8168.00(10.255.105.137), Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  To: 10.255.105.141, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
10.255.105.141                   Rtr   1045   2     2 OSPF(0.0.0.0)
  To: 0102.5502.4211.00(10.255.105.139), Local: 8.42.1.102, Remote: 8.42.1.104

  Local interface index: 0, Remote interface index: 0
  To: 8.31.1.101-1, Local: 8.31.1.103, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0

```

Meaning The routes are appearing in the traffic engineering database.

Related Documentation

- [Link-State Distribution Using BGP Overview on page 8](#)

Dynamic Bandwidth Management Using Container LSP Overview

RSVP LSPs with the autobandwidth feature are increasingly deployed in networks to meet traffic engineering needs. However, the current traffic engineering solutions for point-to-point LSPs are inefficient in terms of network bandwidth utilization, mainly because the ingress routers originating the RSVP LSPs either try to fit the LSPs along a particular path without creating parallel LSPs, or do not interact with the other routers in the network and probe for additional available bandwidth.

This feature provides an ingress router with the capability of acquiring as much network bandwidth as possible by creating parallel LSPs dynamically.

- [Understanding RSVP Multipath Extensions on page 382](#)
- [Junos OS RSVP Multipath Implementation on page 383](#)
- [Current Traffic Engineering Challenges on page 383](#)
- [Using Container LSP as a Solution on page 386](#)
- [Junos OS Container LSP Implementation on page 388](#)
- [Configuration Statements Supported for Container LSPs on page 403](#)
- [Impact of Configuring Container LSPs on Network Performance on page 407](#)
- [Supported and Unsupported Features on page 408](#)

Understanding RSVP Multipath Extensions

The RSVP multipath extensions proposed in the IETF [KOMPELLA-MLSP] allow the setup of traffic engineered multipath label-switched paths (container LSPs). The container LSPs, in addition to conforming to traffic engineering constraints, use multiple independent paths from a source to a destination, thereby facilitating load balancing of traffic. The multipath extensions require changes to the RSVP-TE protocol and allow for merging of labels at the downstream nodes (similar to LDP), which also helps in preserving forwarding resources.

The multipath extensions to RSVP provide the following benefits:

- Ease of configuration. Typically, multiple RSVP LSPs are configured for either load balancing or bin packing. With a container LSP, there is a single entity to provision, manage, and monitor LSPs. Changes in topology are handled easily and autonomously by the ingress LSP, by adding, changing, or removing member LSPs to rebalance traffic, while maintaining the same traffic engineering constraints.
- RSVP equal-cost multipath (ECMP) inherits the standard benefits of ECMP by absorbing traffic surges.
- Multipath traffic engineering allows for better and complete usage of network resources.
- Knowing the relationship among LSPs helps in computing diverse paths with constraint-based routing. It allows adjustment of member LSPs while other member LSPs continue to carry traffic.

- The intermediate routers have an opportunity to merge the labels of member LSPs. This reduces the number of labels that need to get added to the forwarding plane and in turn reduces the convergence time.

If the number of independent ECMP paths is huge, label merging overcomes the platform limitations on maximum (ECMP) next hops. With point-to-point RSVP LSPs that require link or node protection, the next hops are doubled as each LSP is programmed with both primary and backup next hops. RSVP multipath (or ECMP) obviates the need for backup next hops.

- When there is a link failure, the router upstream to the link failure can distribute traffic from the failed link to the remaining ECMP branches, obviating the need for bypass LSPs. The bypass LSP approach not only requires more state when signaling backup LSPs, but also suffers from scaling issues that result in merge-point timing out a protected path state block (PSB) before point of local repair (PLR) gets a chance to signal the backup LSP.

Junos OS RSVP Multipath Implementation

In order to deploy RSVP multipath (ECMP) in a network, all the nodes through which ECMP LSPs pass must understand RSVP ECMP protocol extensions. This can be a challenge, especially in a multivendor networks.


Junos OS implements the RSVP multipath extensions without the need for protocol extensions. A single container LSP, which has the characteristics of ECMP and RSVP TE, is provisioned. A container LSP consists of several member LSPs and is set up between the ingress and egress routing device. Each member LSP takes a different path to the same destination. The ingress routing device is configured with all the required parameters to compute the RSVP ECMP LSP. The parameters configured to compute a set of RSVP point-to-point LSPs can be used by the ingress routing device to compute the container LSP as well.

Current Traffic Engineering Challenges

The main challenge for traffic engineering is to cope with the dynamics of both topology and traffic demands. Mechanisms are needed that can handle traffic load dynamics in scenarios with sudden changes in traffic demand and dynamically distribute traffic to benefit from available resources.

[Figure 36 on page 384](#) illustrates a sample network topology with all the LSPs having the same hold and setup priorities, and admission control restricted on the ingress router. All the links are annotated with a tuple (cost and capacity).

Figure 36: Sample Topology



Some of the traffic engineering problems seen in [Figure 36 on page 384](#) are listed here:

- **Bin Packing**

This problem arises because of a particular order in which LSPs are signaled. The ingress routers might not be able to signal some LSPs with required demands although bandwidth is available in the network, leading to under-utilization of link capacity.

For example, the following LSPs arrive in the sequence mentioned in [Table 7 on page 384](#).

Table 7: LSP Sequence Order for Bin Packing

Time	Source	Destination	Demand	ERO
1	A	E	5	A-C-D-E
2	B	E	10	No ERO

The LSP originating at Router B is not routable as constraint-based routing fails to find a feasible path. However, if Router B is signaled first, both the LSPs are routable. Bin packing happens because of lack of visibility of individual per-LSP, per-device bandwidth demands at the ingress routing device.

Bin packing can also happen when there is no requirement for ordering of LSPs. For example, if there is an LSP with demand X and there are two different paths to the destination from the ingress router with available bandwidths Y1 and Y2, such that Y1 is less than X, Y2 is less than X, and Y1 plus Y2 is greater than or equal to X.

In this case, even though there are enough network resources in terms of available bandwidth to satisfy the aggregate LSP demand X, the LSP might not be signaled or re-optimized with the new demand. In [Figure 36 on page 384](#), with container LSP support, the ingress B creates two LSPs each of size 5 when demand 10 is posed. One LSP is routed along B-C-E and another one along B-C-D-E.

- **Deadlock**

Considering [Figure 36 on page 384](#), the LSPs follow the sequence mentioned in [Table 8 on page 385](#).

Table 8: LSP Sequence Order for Deadlock

Time	Source	Destination	Demand	ERO	Event
1	A	E	2	A-C-D-E	Constraint-based routing with RSVP signaling
2	B	E	2	B-C-D-E	Constraint-based routing with RSVP signaling
3	A	E	2 to 20	A-C-D-E	Constraint-based routing fails, no RSVP signaling

At time 3, the demand on LSP from A to E increases from 2 to 20. If autobandwidth is configured, the change does not get detected until the adjustment timer expires. In the absence of admission control at A, the increased traffic demand might cause traffic to drop on other LSPs that share common links with the mis-behaving LSP.

This happens due to the following reasons:

- Lack of global state at all the ingress routers
- Signaling of mis-behaving demands
- Tearing down of mis-behaving demands

With container LSP configured, ingress A has more chances of splitting the load (even incrementally if not fully) across multiple LSPs. So, LSP from A is less likely to see prolonged traffic loss.

• Latency Inflation

Latency inflation is caused by the autobandwidth and other LSPs parameters. Some of the other factors that contribute to latency inflation include:

- LSP priority

LSPs choose longer paths because shorter paths between data centers located in the same city can be congested. The bandwidth on the shorter paths can get exhausted by equal or higher priority LSPs. Due to periodic LSP optimization by autobandwidth, LSP can get rerouted to a higher delay path. When many LSPs undergo less than optimal path selection, they can potentially form a chain of dependencies. Modifying the LSP priorities dynamically is a workaround to the issue; however, dynamically adjusting LSP priorities to find shorter paths is a challenging task.

- All or Nothing policy

When the demand on an LSP increases and at least one of the links along the shorter path is close to its reservation limit, LSP optimization can force the LSP to move to a longer latency path. LSP has to traverse a long path even though the short path is capable of carrying most of the traffic.

- Minimum and maximum bandwidth

Minimum and maximum bandwidth specify the boundaries for LSP sizes. If minimum bandwidth is small, an LSP is more prone to autobandwidth adjustment because a

small change in bandwidth is enough to cross the threshold limits. LSPs might reroute although bandwidth is available. On the other hand, if the minimum bandwidth is large, network bandwidth might be wasted. If the maximum bandwidth value is small, a large number of LSPs might be needed at the ingress router to accommodate the application demand. If the maximum bandwidth is large, the LSPs can grow larger in size. Such LSPs can suffer because of an all or nothing policy.

- Autobandwidth adjustment threshold

Bandwidth threshold dictates if LSPs need to be re-optimized and resized. If the value is small, LSPs are frequently re-optimized and rerouted. That might cause CPU spike because applications or protocols, such as BGP resolving over the LSPs, might keep the Routing Engine busy doing next-hop resolution. A large value might make an LSP immobile. With container LSP configured, an LSP is less likely to get subjected to one or no policy. An ingress router originates multiple LSPs, although not all LSPs potentially traverse high latency paths.

- Predictability

Service providers often want predictable behavior in terms of how LSPs get signaled and routed. Currently, without any global coordination, it is difficult to set up the same set of LSPs in a predictable way. Consider the two different orderings in [Table 9 on page 386](#) and [Table 10 on page 386](#). The ERO that an LSP uses depends on its signaling time.

Table 9: LSP Sequence Order for Predictability

Time	Source	Destination	Demand	ERO
1	A	E	5	A-C-D-E
2	B	E	5	B-C-E

Table 10: LSP Sequence Order for Predictability

Time	Source	Destination	Demand	ERO
1	B	E	5	B-C-E
2	A	E	5	A-C-D-E

Container LSP does not directly help LSPs find predictable EROs. If LSPs are getting rerouted because of an all or no policy without container LSP configured, such LSPs might see less churn if container LSPs are configured, because smaller LSPs have better chances of finding a shorter or same path.

Using Container LSP as a Solution

A container LSP can be used as a solution to the challenges faced by the current traffic engineering features. Considering [Figure 36 on page 384](#), when the demand X on a container

LSP increases with the network capacity (max-flow) being more than the demand, the following approaches come into effect with a container LSP:

- [Accommodating the New Demand X on page 387](#)
- [Creating New LSPs to Meet Demand X on page 387](#)
- [Assigning Bandwidth to the New LSPs on page 387](#)
- [Controlling the LSP Paths on page 387](#)

Accommodating the New Demand X

In the current implementation, autobandwidth attempts to re-signal an LSP with the new demand X and follows the all or nothing policy as mentioned earlier.

The container LSP approach computes several small (smaller than demand X) bandwidth LSPs such that the aggregate bandwidth is not less than X, and the ingress router performs this adjustment periodically. One of the triggers to create new LSPs or to delete old LSPs can be changed in aggregate bandwidth. The ingress router then load-balances the incoming traffic across the newly created LSPs.

Creating New LSPs to Meet Demand X

Although the number of new LSPs created can be a maximum of the allowed configurable limit, there is not much benefit from these LSPs once the number of LSPs exceeds the number of possible diverse paths or equal-cost multipaths (ECMPs). The benefit of creating the smaller LSPs is seen when an ingress router uses the newly created LSPs for load-balancing traffic. This, however, depends on the network topology and state.

Creating multiple parallel LSPs by all the ingress routers in the network can lead to scaling issues at the transit routers. Thus, the number of new LSPs to be created depends on the size of the individual LSPs and the given aggregate demand, X in this case.

Assigning Bandwidth to the New LSPs

In general, there can be a number of heuristics to allocate bandwidths to the newly created LSPs. An ingress router can solve an optimization problem in which it can maximize a given utility function. The output of an optimization problem is assigning optimal bandwidth values. However, to solve an optimization problem, the number of newly created LSPs has to be fixed. Therefore, it is complex to optimize the number and size of each LSP. Thus, to simplify the problem, the same amount of bandwidth is assumed for all the newly created LSPs, and then the number of required LSPs is computed.

Controlling the LSP Paths

The flexibility to control the LSP paths is expressed in terms of the configuration for point-to-point LSPs and container LSPs. Controlling the LSP paths using the configuration parameters can be applied under two different aspects:

- **Topology**—There are no topology constraints with this feature. Each member LSP is treated like a point-to-point LSP and is re-optimized individually. An ingress router does not try to compute equal IGP cost paths for all its LSPs, but instead it computes paths for all the LSPs using current traffic engineering database information. While computing a path, constraint-based routing adheres to any constraints specified

through the configuration, although there is no change in the constraint-based routing method for path computation.

- **When to create a new LSP**—When to create a new LSP can be explicitly specified. By default, an ingress router periodically computes the aggregate traffic rate by adding up the traffic rate of all the individual LSPs. Looking at the aggregate bandwidth and configuration, the ingress router recomputes the number of LSPs and the bandwidths of the LSPs. The new LSPs are then signaled or the existing LSPs are re-signaled with the updated bandwidth. Instead of looking at the instantaneous aggregate rate, the ingress routers can compute an average (of aggregates) over some duration by removing outlier samples (of aggregates). Managing the LSPs that remain outstanding and active by considering aggregate bandwidth is more scalable than creating the new LSPs based on the usage of a particular LSP. The intervals and thresholds can be configured to track the aggregate traffic and trigger adjustment. These dynamic LSPs co-exist and interoperate with per-LSP autobandwidth configuration.

Junos OS Container LSP Implementation

A container LSP is an ECMP TE LSP that acts like a container LSP consisting of one or more member LSPs. A point-to-point TE LSP is equivalent to a container LSP with a single member LSP. Member LSPs are added to the container LSP through a process called splitting, and removed from the container LSP through a process called merging.

- [Container LSP Terminology on page 388](#)
- [LSP Splitting on page 389](#)
- [LSP Merging on page 391](#)
- [Node and Link Protection on page 393](#)
- [Naming Convention on page 393](#)
- [Normalization on page 394](#)
- [Constraint-Based Routing Path Computation on page 399](#)
- [Sampling on page 400](#)
- [Support for NSR, IPG-FA, and Static Routes on page 400](#)

Container LSP Terminology

The following terms are defined in the context of a container LSP:

- **Normalization**—An event occurring periodically when an action is taken to adjust the member LSPs, either to adjust their bandwidths, their number, or both. A normalization process is associated with a sampling process and periodically estimates aggregate utilization of a container LSP.
- **Nominal LSP**—The instance of a container LSP that is always present.
- **Supplementary LSP**—The instances or sub-LSPs of a container LSP, which are dynamically created or removed.

Autobandwidth is run over each of the member LSPs, and each LSP is resized according to the traffic it carries and the autobandwidth configuration parameters. The aggregate

demand on a container LSP is tracked by adding up the bandwidth across all the member LSPs.

- **Minimum signaling-bandwidth**—The minimum bandwidth with which a member LSP is signaled at the time of normalization or initialization. This could be different from the minimum-bandwidth defined under autobandwidth.
- **Maximum signaling-bandwidth** —The maximum bandwidth with which a member LSP is signaled at the time of normalization or initialization. This could be different from the maximum-bandwidth defined under autobandwidth.
- **Merging-bandwidth** —Specifies the lower bandwidth threshold on the aggregate bandwidth usage, such that if the aggregate usage falls below this value, the ingress router merges the member LSPs at the time of normalization.
- **Splitting-bandwidth** —Specifies the upper bandwidth threshold on the aggregate bandwidth usage, such that if the aggregate usage exceeds this value, the ingress router splits the member LSPs at the time of normalization.
- **Aggregate minimum-bandwidth** —Sum of merging-bandwidth of the current active member LSPs. This minimum bandwidth is different from the autobandwidth minimum-bandwidth.
- **Aggregate maximum-bandwidth**—Sum of the splitting-bandwidth of the current active member LSPs. This maximum bandwidth is different from the autobandwidth maximum-bandwidth.

LSP Splitting

- [Operational Overview on page 389](#)
- [Operational Constraints on page 390](#)
- [Supported Criteria on page 390](#)
- [Splitting Triggers on page 391](#)

Operational Overview

The LSP splitting mechanism enables an ingress router to create new member LSPs or to re-signal existing LSPs with different bandwidths within a container LSP when a demand X is placed on the container LSP. With LSP splitting enabled, an ingress router periodically creates a number of LSPs (by signaling new ones or re-signaling existing ones) to accommodate a new aggregate demand X. In the current implementation, an ingress router tries to find an LSP path satisfying a demand X and other constraints. If no path is found, either the LSP is not signaled or it remains up, but with the old reserved bandwidth.

Between two normalization events (splitting or merging), individual LSPs might get re-sigaled with different bandwidths due to the autobandwidth adjustments. If a container LSP is not configured with autobandwidth, then the member LSPs are signaled with the static bandwidth value, if configured. There is no dynamic splitting in this case, as there is no dynamic estimation of aggregate bandwidth. The splitting adjustments with a specific bandwidth value can be manually triggered.

**NOTE:**

Be aware of the following considerations for LSP splitting:

- After LSP splitting, the ingress router continues to inject one forwarding adjacency. Forwarding adjacencies are not supported in IGP for this feature.
- Between two normalization events, two LSPs might have different bandwidths subjected to autobandwidth constraints.
- After LSPs are split (or merged), make-before-break uses the fixed filter (FF) style sharing unless the adaptive option is configured. However, two different LSPs do not do the shared explicit (SE) style sharing for this feature.
- When LSPs are re-signaled with modified bandwidths, some of the LSPs might not get signaled successfully, leading to failover options.

Operational Constraints

LSP splitting has the following operational constraints:

- LSP bandwidth—Although there are a number of ways to allocate bandwidth values to the LSPs, the Junos OS implementation supports only an equal-bandwidth allocation policy when normalization is done, wherein all the member LSPs are signaled or re-signaled with equal bandwidth.
- Number of LSPs—If an ingress router is configured to have a minimum number of LSPs, it maintains the minimum number of LSPs even if the demand can be satisfied with less than the minimum number of LSPs. In case the ingress router is unable to do constraint-based routing for computations on the sufficient number of LSPs or signal sufficient number of LSPs, the ingress router resorts to a number of fallback options.

By default, an incremental approach is supported as a fallback option (unless configured differently), where an ingress router makes attempts to bring up the sufficient number of LSPs, such that the new aggregate bandwidth exceeds the old aggregate bandwidth (and is as close to the desired demand as possible). The ingress router then load-balances traffic using the LSPs. The LSPs that could not be brought up are removed by the ingress router.

Supported Criteria

When a container LSP signals a member LSP, the member LSP gets signaled with minimum-signaling-bandwidth. Since each member LSP is configured with autobandwidth, between two normalization events, each LSP can undergo autobandwidth adjustment multiple times. As the traffic demand increases, the ingress router creates additional supplementary LSPs. All member LSPs are used for ECMP, so they should roughly have the same reserved bandwidth after normalization.

For example, if there are K LSPs signaled after normalization, each LSP is signaled with equal bandwidth B. The total aggregate bandwidth reserved is B.K, where B satisfies the following condition:

- Minimum signaling-bandwidth is less than or equal to B, which in turn is less than or equal to the maximum signaling-bandwidth

(minimum-signaling-bandwidth \leq B \leq maximum-signaling-bandwidth)

Until the next normalization event, each member LSP undergoes several autobandwidth adjustments. After any autobandwidth adjustment, if there are N LSPs with reserved bandwidths b_i , where $i=1,2,\dots, N$, each b_i should satisfy the following condition:

- Minimum bandwidth is less than or equal to b_i , which in turn is less than or equal to the maximum bandwidth

(minimum-bandwidth $\leq b_i \leq$ maximum-bandwidth)

Both the above-mentioned conditions are applicable for per member LSP (nominal and supplementary), and essentially have the reserved bandwidth to exist within a range.

Splitting Triggers

Every time the normalization timer expires, the ingress router decides if LSP splitting is required. The ingress router works with the aggregate bandwidth instead of the individual LSP bandwidths. The following two variables are defined for aggregate bandwidth:

- **Current-Aggr-Bw**—Sum of reserved bandwidths of all current member LSPs.
- **New-Aggr-Bw**—Sum of traffic rates on all current member LSPs based on sampling.

Taking for example, if there are N member LSPs in the network at the time of normalization, the two approaches to trigger LSP splitting are as follows:

- **Absolute trigger**—LSP splitting is performed when **New-Aggr-Bw** is greater than **Aggregate-maximum-bandwidth**.

(**New-Aggr-Bw** > **Aggregate-maximum-bandwidth**)

- **Relative trigger**—The **Current-Aggr-Bw** is compared with **New-Aggr-Bw** at the ingress routing device. LSP splitting is performed when the difference in the bandwidth amount is off by a threshold.

($[1-a] \times \text{Current-Aggr-Bw} < \text{New-Aggr-Bw} < [1+a] \times \text{Current-Aggr-Bw}$, where $0 \leq a \leq 1$)

When **New-Aggr-Bw** is greater than or equal to $[1+a]$ multiplied by **Current-Aggr-Bw**, the ingress routing device does not perform normalization, but instead LSP splitting is done. However, when both LSP splitting and LSP merging are configured on the ingress router, LSP splitting is triggered on the ingress router when one of the two conditions is satisfied.

LSP Merging

- [Operational Overview on page 392](#)
- [Operational Constraints on page 392](#)
- [Merging Triggers on page 392](#)

Operational Overview

Junos OS supports two kinds of LSPs – CLI-configured LSPs and dynamically created LSPs. The CLI-configured LSPs are created manually and remain in the system until the configuration is modified. The dynamic LSPs are created dynamically by next generation MVPN, BGP virtual private LAN service (VPLS), or LDP, based on a template configuration, and are removed from the system when not used by any application for a certain duration. LSP merging follows a similar approach as dynamic LSPs.

LSP merging enables an ingress routing device to dynamically eliminate some member LSPs of the container LSP so less state information is maintained in the network. If an ingress router provisions several member LSPs between the ingress and egress routers, and there is an overall reduction in aggregate bandwidth (resulting in some LSPs being under-utilized), the ingress router distributes the new traffic load among fewer LSPs.

Although there are a number of ways to merge the member LSPs, Junos OS supports only overall-merge when normalization is being performed. An ingress router considers the aggregate demand and the minimum (or maximum) number of LSPs and revises the number of LSPs that should be active at an ingress routing device. As a result, the following can take place periodically as the normalization timer fires:

- Re-signaling some of the existing LSPs with updated bandwidth
- Creating new LSPs
- Removing some of the existing LSPs

Operational Constraints

If a container LSP is not configured with autobandwidth, then the member LSPs are signaled with the static bandwidth value, if configured. LSP merging does not happen because there is no dynamic estimation of aggregate bandwidth. However, a manual trigger for splitting and adjusting with a specific bandwidth value can be configured.



NOTE:

- Nominal LSPs are never deleted as part of LSP merging.
 - Before deleting an LSP, the LSP is made inactive, so that traffic shifts to other LSPs before removing the LSP. This is because RSVP sends PathTear before deleting routes and next hops from the Packet Forwarding Engine.
 - When member LSPs are re-signaled with modified bandwidth, it might happen that some LSPs do not get signaled successfully.
-

Merging Triggers

Every time the normalization timer expires, the ingress router decides if LSP merging is required. The ingress router works with the aggregate bandwidth instead of the individual LSP bandwidths. The following two variables are defined for aggregate bandwidth:

- **Current-Aggr-Bw**—Sum of reserved bandwidths of all current member LSPs.

- **New-Aggr-Bw**—Sum of traffic rates on all current member LSPs based on sampling.

For example, if there are N member LSPs in the network at the time of normalization, the two approaches to trigger LSP merging are as follows:

- Absolute trigger—LSP merging is performed when **New-Aggr-Bw** is less than **Aggregate-minimum-bandwidth**.

$$(\text{New-Aggr-Bw} < \text{Aggregate-maximum-bandwidth})$$

- Relative trigger—The **Current-Aggr-Bw** is compared with **New-Aggr-Bw** at the ingress routing device. LSP merging is performed when the difference in the bandwidth amount is off by a threshold.

$$([1-a] \times \text{Current-Aggr-Bw} < \text{New-Aggr-Bw} < [1+a] \times \text{Current-Aggr-Bw}, \text{ where } 0 \leq a \leq 1)$$

When the **New-Aggr-Bw** value is less than or equal to $[1+a]$ multiplied by the **Current-Aggr-Bw** value, the ingress routing device does not perform normalization, but instead LSP merging is done. However, when both LSP splitting and LSP merging are configured on the ingress router, LSP splitting is triggered on the ingress router when one of the two conditions is satisfied.

Node and Link Protection

Junos OS supports the following mechanisms for node and link protection:

- Fast-reroute
- Link protection
- Node-link protection

Only one of the above-mentioned modes of protection can be configured on an ingress routing device at any given time. All member LSPs (nominal and supplementary) use the same mode of protection that is configured.

Naming Convention

While configuring a container LSP, a name is assigned to the LSP. The name of a nominal and a supplementary LSP is formed by adding the configured-name suffix and an auto-generated suffix to the name of the container LSP. The name of the container LSP is unique and is checked for accuracy during the configuration parsing. The container LSP name should uniquely identify parameters, such as the ingress and egress router names.



NOTE: A container LSP member LSP and a point-to-point LSP on an ingress routing device cannot have the same LSP name.

The container LSPs follow a number-based LSP naming convention. For example, if the nominal LSP's configured name is **bob** and the number of member LSPs is N, the member LSPs are named **bob-*<configured-suffix>-1***, **bob-*<configured-suffix>-2***, ..., and **bob-*<configured-suffix>-N***.

After a normalization event, the number of member LSPs can change. For example, if the number of member LSPs increases from six to eight, then the ingress routing device keeps the first six LSPs named **bob-*<configured-suffix>-1***, **bob-*<configured-suffix>-2***, ..., and **bob-*<configured-suffix>-6***. The two additional LSPs are named **bob-7** and **bob-8**. The original LSPs might need to be re-optimized if their signaled bandwidth changes.

Similarly, if the number of member LSPs reduces from eight to six, the ingress routing device re-signals the member LSPs in such a way that the remaining active LSPs in the system are named **bob-*<configured-suffix>-1***, **bob-*<configured-suffix>-2***, ..., and **bob-*<configured-suffix>-6***.

In the process of creating new LSPs, an RSVP LSP named **bob-*<configured-suffix>-7*** can be configured.

Normalization

- [Operational Overview on page 394](#)
- [Operational Constraints on page 394](#)
- [Inter-Operation with Autobandwidth on page 395](#)

Operational Overview

Normalization is an event that happens periodically. When it happens, a decision is made on the number of member LSPs that should remain active and their respective bandwidths in a container LSP. More specifically, the decision is made on whether new supplementary LSPs are to be created, or any existing LSPs are required to be re-signaled or deleted during the normalization event.

Between two normalization events, a member LSP can undergo several autobandwidth adjustments. A normalization timer, similar to re-optimization timer, is configured. The normalization timer interval should be no less than the adjustment interval or optimization timer.



NOTE: Normalization is not triggered based on network events, such as topology changes.

Operational Constraints

Normalization has the following operational constraints:

- Normalization happens only when none of the member LSPs are undergoing re-optimization or make-before-break. Normalization starts when all the member LSPs complete their ongoing make-before-break. If normalization is pending, new optimization should not be attempted until the normalization is complete.
- After normalization, an ingress routing device first computes a set of bandwidth-feasible paths using constraint-based routing computations. If enough constraint-based routing computed paths are not brought up with an aggregate bandwidth value that exceeds the desired bandwidth, several failover actions are taken.

- After a set of bandwidth-feasible paths are available, the ingress routing device signals those paths while keeping the original set of paths up with the old bandwidth values. The make-before-break is done with shared explicit (SE) sharing style, and when some of the LSPs do not get successfully re-signaled, a bounded number of retries is attempted for a specified duration. Only when all the LSPs are successfully signaled does the ingress router switch from the old instance of the container LSP to the newer instance. If all LSPs could not be successfully signaled, the ingress router keeps those instances of members that are up with higher bandwidth values.

For example, if the bandwidth of an old instance of a member LSP (LSP-1) is 1G, the LSP is split into LSP-1 with bandwidth 2G and LSP-2 with bandwidth 2G. If the signaling of LSP-1 with bandwidth 2G fails, the ingress router keeps LSP-1 with bandwidth 1G and LSP-2 with bandwidth 2G.

When there is a signaling failure, the ingress routing device stays in the error state, where some LSPs have updated bandwidth values only if the aggregate bandwidth has increased. The ingress router makes an attempt to bring up those LSPs that could not be successfully signaled, resulting in minimum traffic loss.

- If an LSP goes down between two normalization events, it can increase the load on other LSPs that are up. In order to prevent overuse of other LSPs, premature normalization can be configured in case of LSP failure. LSPs can go down because of pre-emption or lack of node or link protection. It might not be necessary to bring up the LSPs that are down because the normalization process re-runs the constraint-based routing path computations.

Inter-Operation with Autobandwidth

Taking as an example, there is one nominal LSP named LSP-1 configured with the following parameters:

- Splitting-bandwidth and maximum-signaling-bandwidth of 1G
- Merging-bandwidth and minimum-signaling-bandwidth of 0.8G
- Autobandwidth

Normalization is performed differently in the following scenarios:

- [Changes in Per-LSP Autobandwidth Adjustments on page 395](#)
- [Changes in Traffic Growth on page 397](#)
- [Computed Range and Configured Feasible Ranges on page 397](#)

Changes in Per-LSP Autobandwidth Adjustments

[Table 11 on page 396](#) illustrates how normalization splits and merges member LSPs as autobandwidth adjustments change per-LSP bandwidth with unconditional normalization.

Table 11: Normalization with Per-LSP Autobandwidth Adjustment Changes

Normalization Time	Current State	Events	Adjusted State
T0	No state.	Initialization	LSP-1 is signaled with bandwidth of 0.8G
T1	LSP-1 usage increases to 1.5G	<ul style="list-style-type: none"> Multiple autobandwidth adjustments since T0 is possible. The ingress router decides to split LSP-1 into two LSPs, and creates LSP-2. 	LSP-1 = 0.8G LSP-2 = 0.8G
T2	LSP-1 usage increase to 2G LSP-2 usage increases to 0.9G (within limits)	<ul style="list-style-type: none"> Aggregate bandwidth is 2.9G, which exceeds aggregate splitting maximum of 2G. The ingress router decides to split LSP-1 into three LSPs, and creates LSP-3. 	LSP-1 = 1G LSP-2 = 1G LSP-3 = 1G
T3	LSP-3 usage increases to 1.5G	<ul style="list-style-type: none"> Aggregate bandwidth is 3.5G with a maximum aggregate splitting of 3G. The ingress router decides to split LSP-1 into four LSPs, and creates LSP-4. 	LSP-1 = 1G LSP-2 = 1G LSP-3 = 1G LSP-4 = 1G
T4	LSP-2 usage drops to 0.5G	<ul style="list-style-type: none"> Aggregate bandwidth is 3G. The ingress router decides to merge LSP-1 and removes LSP-4. 	LSP-1 = 1G LSP-2 = 1G LSP-3 = 1G

Because autobandwidth is configured on a per-LSP basis, every time there is an autobandwidth adjustment, the ingress router re-signals each LSP with **Max Avg Bw**.

Another approach to handling the changes in per-LSP autobandwidth adjustments is to not allow individual LSPs to run autobandwidth on the ingress router, but to run autobandwidth in passive (monitor) mode. This way, sampling is done at every statistics interval for member LSPs only, and normalization is performed for the container LSP alone instead of acting on individual LSPs adjustment timer expiry.

As a result, the number of re-signaling attempts and bandwidth fluctuations for a given member LSP is reduced. Only the computed bandwidth-values per-member LSP is used by the ingress router to find an aggregate bandwidth to be used during normalization. Configuring autobandwidth adjustment followed by normalization (adjustments and normalization intervals are comparable) can lead to considerable overhead because of re-signaling.

Taking the same example, and applying the second approach, LSP-1 goes from 0.8G to 1.5G and then back to 0.8G. If the normalization timer is of the same order as the adjustment interval, the ingress router leaves LSP-1 alone with its original 0.8G and only signals LSP-2 with 0.8G. This helps achieve the final result of normalization, thus avoiding the extra signaling attempt on LSP-1 with 1.5G at adjustment timer expiry.

Because member LSPs always use equal bandwidth, any adjustment done on member LSPs is undone. The member LSPs are re-signaled with reduced bandwidth when compared to the reserved capacity in adjustment trigger with normalization trigger. Therefore, avoiding adjustment trigger for member LSPs might be useful assuming that normalization and adjustment intervals are of the same order.



NOTE: We recommend that the normalization timer be higher than the autobandwidth adjustment interval and regular optimization duration, as the traffic trends are observed at a longer time scale and normalization is performed one-to-three times per day. An LSP can undergo optimization for the following reasons:

- Normal optimization
- Autobandwidth adjustment
- Normalization

Changes in Traffic Growth

Table 12 on page 397 illustrates how normalization is performed when traffic grows in large factor.

Table 12: Normalization with Traffic Growth

Normalization Time	Current State	Events	Adjusted State
T0	No state		LSP-1 is signaled with bandwidth of 0.8G
T1	LSP-1 usage increase to 3G	<ul style="list-style-type: none"> • Aggregate usage exceeds maximum splitting bandwidth • The ingress router decides to split LSP-1, and creates two more supplementary LSPs 	LSP-1 = 1G LSP-2 = 1G LSP-3 = 1G

Having fewer LSPs is preferred over signaling four LSPs each with 0.8G bandwidth, unless there is a constraint on the minimum number of LSPs.

Computed Range and Configured Feasible Ranges

When an ingress router is configured with the minimum and maximum number of LSPs, and per LSP splitting-bandwidth and merging-bandwidth values, the bandwidth thresholds are used for splitting and merging. For this, the number of LSPs (N) should satisfy the following constraints:

$$\text{minimum-member-lsps} \leq N \leq \text{maximum-member-lsps}$$

At the time of normalization, based on the aggregate demand X:

$$\lceil X/\text{splitting-bandwidth} \rceil \leq N \leq \lfloor X/\text{merging-bandwidth} \rfloor$$

The above-mentioned constraints provide two ranges for N to work from. If the two ranges for N are overlapping, N will be selected from the overlapping interval (lowest possible N) to keep the number of LSPs small in the network.

Otherwise, if maximum-member-lsps is less than $\lceil X/\text{splitting-bandwidth} \rceil$, the ingress router keeps (at maximum) the maximum-member-lsps in the system, and the bandwidth of each LSP is $\lceil X/\text{maximum-member-lsps} \rceil$ or the maximum-signaling-bandwidth, whichever is less. It is possible that some LSPs might not get signaled successfully.

Similarly, if minimum-member-lsps is greater than $\lceil X/\text{merging-bandwidth} \rceil$, the ingress router keeps (at minimum) the minimum-member-lsps in the system, and the bandwidth of each LSP is $\lceil X/\text{minimum-member-lsps} \rceil$ or the minimum-signaling-bandwidth, whichever is less.

Taking as an example, normalization is performed as following in these cases:

- Case 1
 - minimum-member-lsps = 2
 - maximum-member-lsps = 10
 - aggregate demand = 10G
 - merging-bandwidth = 1G
 - splitting-bandwidth = 2.5G

In this case, the ingress routing device signals four member LSPs each with a bandwidth of 2G.

- Case 2
 - minimum-member-lsps = 5
 - maximum-member-lsps = 10
 - aggregate demand = 10G
 - merging-bandwidth = 2.5G
 - splitting-bandwidth = 10G

In this case, the ingress routing device signals five member LSPs each with a bandwidth of 2G. Here, the static configuration on the number of member LSPs takes precedence.

- Case 3
 - minimum-signaling-bandwidth = 5G
 - maximum-signaling-bandwidth = 40G
 - merging-bandwidth = 10G
 - splitting-bandwidth = 50G

When a container LSP comes up, the nominal LSP is signaled with minimum-signaling-bandwidth. At the time of normalization, the

new-aggregate-bandwidth is 100G. To find N and the bandwidth of each LSP, N should satisfy the following constraint:

$$100/50 \leq N \leq 100/10, \text{ which gives } 2 \leq N \leq 10$$

Therefore, N is equal to:

- N = 2, bandwidth = $\min \{100/2G, 40G\} = 40G$

This option does not satisfy the new aggregate of 100G.

- N = 3, bandwidth = $\min \{100/3G, 40G\} = 33.3G$

This option makes the aggregate bandwidth equal to 100G.

In this case, the ingress routing device signals three LSPs each with a bandwidth of 33.3G.



NOTE: The ingress router does not signal an LSP smaller than the minimum-signaling-bandwidth.

Constraint-Based Routing Path Computation

Although there are no changes in the general constraint-based routing path computation, with a container LSP, there is a separate module that oversees the normalization process, schedules constraint-based routing events, and schedules switchover from an old instance to a new instance, when appropriate. An ingress routing device has to handle the constraint-based routing path computation periodically. When normalization occurs, an ingress router has to compute constraint-based routing paths, if the number of LSPs or the bandwidth of the LSPs needs to be changed.

For example, there are K LSPs at the ingress router with bandwidth values X-1, X-2, ..., and X-K. The current aggregate bandwidth value is Y, which is the sum of X-1 plus X-2 plus X-K. If there is a new demand of W, the ingress router first computes how many LSPs are required. If the ingress router only needs N LSPs (LSP-1, LSP-2, ..., and LSP-N) each with bandwidth value B, the task of the constraint-based routing module is to provide a set of bandwidth-feasible LSPs that can accommodate the new aggregate demand which is not less than Y.

The ingress router then tries to see if the constraint-based routing paths can be computed successfully for all N LSPs. If the paths for all the LSPs are found successfully, the constraint-based routing module returns the set to the normalization module.

It is possible that the constraint-based routing computation is not successful for some LSPs. In this case, the ingress routing device takes the following action:

- If the configuration allows for incremental-normalization, implying if the ingress router has enough LSPs whose aggregate exceeds Y, the constraint-based routing module returns that set of paths.
- Whether increment-normalization is configured or not, if constraint-based routing paths could not be computed for a sufficient number of LSPs, the ingress router has to repeat the process of finding a new set of LSPs. Initially, the ingress router starts

with the lowest value of N from the feasible region. Every time, the ingress router has to revise the number, it linearly increases it by 1. As a result, per LSP bandwidth becomes less and therefore, there is a greater chance of successful signaling. The process is repeated for all feasible values of N (or some bounded number of times or duration as configured).

The ingress router signals the LSPs after successful computations of the constraint-based routing path computation. It might happen that when the LSPs are signaled, signaling of many LSPs fail. In addition to the constraint-based routing path computations to be successful, the RSVP signaling should also succeed, such that the new aggregate is not less than the old aggregate bandwidth.

Sampling

Sampling is important for normalization to function. With sampling configured, an ingress routing device is able to make a statistical estimate of the aggregate traffic demands. Every time the sampling timer fires, the ingress routing device can consider traffic rates on different LSPs and compute an aggregate bandwidth sample. This sampling timer is different from the statistics sampling done periodically by RSVP on all LSPs. The aggregate bandwidth is a sample to be used at the time of normalization. An ingress routing device can save past samples to compute an average (or some other statistical measure) and use it the next time normalization happens.

To remove any outlier samples, a sampling token is configured. In other words, from all the aggregate samples collected during the configured time, the bottom and top outliers are ignored before computing a statistical measure from the remaining samples.

The following two methods of computing an aggregate bandwidth value are supported:

- **Average**—All the aggregate bandwidth samples are considered by the ingress routing device, and then all the outlier samples are removed. The average bandwidth value is computed from the remaining samples to be used during normalization.
- **Max**—All the aggregate bandwidth samples are considered by the ingress routing device, and then all the outlier samples are removed. The maximum bandwidth value is picked from the remaining samples to be used during normalization.

The time duration, the number of past aggregate samples to store, the percentile value to determine, and the ignore outliers are user-configurable parameters.

Support for NSR, IPG-FA, and Static Routes

Starting with Junos OS Release 15.1, container label-switched paths (LSPs) provide support for nonstop active routing (NSR), IGP forwarding adjacency (FA), and static routes to address the requirements of wider business cases.

- [NSR Support on page 401](#)
- [IPG-FA Support on page 402](#)
- [Static Route Support on page 403](#)

NSR Support

A container LSP has the characteristics of ECMP and RSVP traffic engineering. Because a container LSP consists of several member LSPs between an ingress and an egress router, with each member LSP taking a different path to the same destination, the ingress router is configured with all the parameters necessary to compute an RSVP ECMP LSP. These parameters along with the forwarding state information have to be synchronized between the master and backup Routing Engines to enable the support for nonstop active routing (NSR) for container LSPs. While some of the forwarding state information on the backup Routing Engine is locally built based on the configuration, most of it is built based on periodic updates from the master Routing Engine. The container LSPs are created dynamically using the replicated states on the backup Routing Engine.

By default, normalization occurs once in every 6 hours and during this time, a number of autobandwidth adjustments happen over each member LSP. A member LSP is resized according to the traffic it carries and the configured autobandwidth configuration parameters. The aggregate demand on a container LSP is tracked by summing up the bandwidth across all the member LSPs.

For RSVP point-to-point LSPs, a Routing Engine switchover can be under any one of the following:

- **Steady state**

In the steady state, the LSP state is up and forwards traffic; however, no other event, such as the make-before-break (MBB), occurs on the LSP. At this stage, the RPD runs on both the Routing Engines, and the switchover event toggles between the master and backup Routing Engine. The backup Routing Engine has the LSP information replicated already. After the switchover, the new master uses the information of the replicated structure to construct the container LSP and en-queues the path (ERO) of LSP in the retrace mode. RSVP signals and checks if the path mentioned in the ERO is reachable. If the RSVP checks fail, then the LSP is restarted. If the RSVP checks succeed, the LSP state remains up.

- **Action leading to make-before-break (MBB)**

A container LSP can be optimized with updated bandwidth, and this change is done in a MBB fashion. During an MBB process, there are two path instances for a given LSP, and the LSP switches from one instance to another. For every Routing Engine switchover, the path is checked to find out where in the MBB process the path is. If the path is in the middle of the MBB process, with the main instance being down and the re-optimized path being up, then MBB can switch over to the new instance. The **show mpls lsp extensive** command output, in this case, is as follows:

```
13 Dec 3 01:33:38.941 Make-before-break: Switched to new instance
12 Dec 3 01:33:37.943 Record Route: 10.1.1.1
11 Dec 3 01:33:37.942 Up
10 Dec 3 01:33:37.942 Automatic Autobw adjustment succeeded: BW changes
from 100 bps to 281669 bps
9 Dec 3 01:33:37.932 Originate make-before-break call
8 Dec 3 01:33:37.931 CSPF: computation result accepted 10.1.1.1
7 Dec 3 01:28:44.228 CSPF: ERO retrace was successful 10.1.1.1
6 Dec 3 01:19:39.931 10.1.1.2 Down: mbb/reopt
5 Dec 3 01:18:29.286 Up: mbb/reopt
```

```

4 Dec 3 01:14:47.119 10.1.1.2 Down: mbb/reopt
3 Dec 3 01:13:29.285 Up: mbb/reopt
2 Dec 3 01:10:59.755 Selected as active path: selected by master RE

```

A similar behavior is retained for member LSPs during bandwidth optimization.

A Routing Engine switchover under the steady state (when normalization is not in progress), keeps the container LSPs up and running without any traffic loss. Events, such as an MBB due to autobandwidth adjustments, link status being down, or double failure, in the steady state are similar to a normal RSVP point-to-point LSP.

If the container LSP is in the process of normalization, and the normalization event is triggered either manually or periodically, it goes through the computation and execution phase. In either of the cases, zero percent traffic loss is not guaranteed.

- Normalization in the computation phase

During the computation phase, the master Routing Engine calculates the targeted member LSP count and bandwidth with which each member LSP should be re-signaled. The backup Routing Engine has limited information about the container LSP, such as the LSP name, LSP ID, current bandwidth of its member LSP, member LSP count, and the normalization retry count. If the switchover happens during the computation phase, then the backup Routing Engine is not aware of the targeted member LSP count and the bandwidth to be signaled. Since traffic statistics are not copied to the backup Routing Engine, it cannot compute the targeted member count and bandwidth. In this case, the new master Routing Engine uses the old data stored in the targeted member LSP count and the targeted bandwidth to signal the LSPs.

- Normalization in the execution phase

During the execution phase, RSVP of the master Routing Engine tries to signal the LSPs with the newly calculated bandwidth. If the switchover occurs during the signaling of LSPs with greater bandwidth or during LSP splitting or merging, then the new master Routing Engine uses the information of the targeted member count and bandwidth value to be signaled with, to bring up the LSPs.

IPG-FA Support

A forwarding adjacency (FA) is a traffic engineering label-switched path (LSP) that is configured between two nodes and used by an interior gateway protocol (IGP) to forward traffic. By default, an IGP does not consider MPLS traffic-engineering tunnels between sites, for traffic forwarding. Forwarding adjacency treats a traffic engineering LSP tunnel as a link in an IGP topology, thus allowing the nodes in the network also to forward the IP traffic to reach the destination over this FA LSP. A forwarding adjacency can be created between routing devices regardless of their location in the network.

To advertise a container LSP as an IGP-FA, the LSP name needs to be configured either under IS-IS or OSPF. For example:

```

IS-IS [edit]
      protocols {
        isis {
          label-switched-path container-lsp-name;
        }
      }

```



```

OSPF    [edit]
        protocols {
          ospf {
            area 0.0.0.0 {
              label-switched-path container-lsp-name;
            }
          }
        }

```



NOTE: The IGP-FA is applied to both container LSPs and regular point-to-point LSPs. If a container LSP and a point-to-point LSP share the same name, the point-to-point LSP is given preference for FA.

Static Route Support

Static routes often include only one or very few paths to a destination and generally do not change. These routes are used for stitching services when policies and other protocols are not configured.

To advertise a container LSP as a static route, the LSP name needs to be configured under the static route configuration. For example:

```

Static Route  [edit]
              routing-options {
                static {
                  route destination {
                    lsp-next-hop container-lsp-name;
                  }
                }
              }

```



NOTE: The static route support is applied to both container LSPs and regular point-to-point LSPs. If a container LSP and a point-to-point LSP share the same name, the point-to-point LSP is given preference for static routing.

Configuration Statements Supported for Container LSPs

Table 13 on page 404 lists the MPLS LSP configuration statements that apply to RSVP LSP and a container LSP (nominal and supplementary).

The configuration support is defined using the following terms:

- Yes—The configuration statement is supported for this type of LSP.
- No—The configuration statement is not supported for this type of LSP.
- N/A—The configuration statement is not applicable for this type of LSP.

Table 13: Applicability of RSVP LSPs Configuration to a Container LSP

Configuration Statement	RSVP LSP (Ingress)	Member LSP (Ingress)
adaptive (Default: non-adaptive)	Yes	Yes
admin-down	Yes	Yes
admin-group	Yes	Yes
admin-groups-except	Yes	Yes
apply-groups	Yes	Yes
apply-groups-except	Yes	Yes
associate-backup-pe-groups	Yes	No
associate-lsp (No bidirectional support)	Yes	No
auto-bandwidth	Yes	Yes
backup	Yes	No
bandwidth	Yes	Yes
class-of-service	Yes	Yes
corouted-bidirectional (No bidirectional support)	Yes	No
corouted-bidirectional-passive (No bidirectional support)	Yes	No
description	Yes	Yes
disable	Yes	Yes
egress-protection	Yes	No
exclude-srlg	Yes	Yes
fast-reroute (Same fast reroute for all member LSPs)	Yes	Yes

Table 13: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

Configuration Statement	RSVP LSP (Ingress)	Member LSP (Ingress)
from	Yes	Yes
hop-limit	Yes	Yes
install	Yes	Yes
inter-domain (Same termination router)	Yes	Yes
secondary (All LSPs are primary)	Yes	No
ldp-tunneling (All LSPs do tunneling)	Yes	Yes
least-fill	Yes	Yes
link-protection (All LSPs share same link protection mechanism)	Yes	Yes
lsp-attributes	Yes	Yes
lsp-external-controller	Yes	No
metric (All LSPs are same)	Yes	Yes
most-fill	Yes	Yes
no-cspf (LSPs use IGP)	Yes	Yes
no-decrement-ttl (All LSPs share same TTL behavior)	Yes	Yes
no-install-to-address	Yes	Yes
no-record	Yes	Yes
node-link-protection (All LSPs share same node-link protection mechanism)	Yes	Yes

Table 13: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

Configuration Statement	RSVP LSP (Ingress)	Member LSP (Ingress)
oam	Yes	Yes
optimize-hold-dead-delay (All LSPs have same value)	Yes	Yes
optimize-switchover-delay (All LSPs have same value)	Yes	Yes
optimize-timer (All LSPs have same value)	Yes	Yes
p2mp	Yes	N/A
policing (Variable traffic)	Yes	No
preference	Yes	Yes
primary (All paths are primary)	Yes	No
random	Yes	Yes
record	Yes	Yes
retry-limit (Applicable to members)	Yes	Yes
retry-timer (Applicable to members)	Yes	Yes
revert-timer (No secondary LSP)	Yes	No
secondary (All LSPs are primary)	Yes	No
soft-preemption	Yes	Yes

Table 13: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

Configuration Statement	RSVP LSP (Ingress)	Member LSP (Ingress)
standby (All LSPs are standby)	Yes	No
template	Yes	No
to	Yes	Yes
traceoptions	Yes	Yes
ultimate-hop-popping	Yes	Yes

Impact of Configuring Container LSPs on Network Performance

A container LSP is a container LSP that allows multiple member LSPs to co-exist and be managed as a bundle. The member LSPs are similar to independent point-to-point RSVP LSPs. As a result, resource consumption is similar to the sum of resources consumed by each point-to-point RSVP LSP. However, provisioning a container LSP is more efficient, as under-utilized member LSPs are dynamically removed, thus saving memory and CPU resources.

The container LSP features are dependent on the presence of a functional base MPLS RSVP implementation. As a result, a container LSP does not introduce any security considerations beyond the existing considerations for the base MPLS RSVP functionality. The categories of possible attacks and countermeasures are as follows:

- Interaction with processes and router configuration

No new communication mechanisms with external hosts are required for a container LSP. Data arrives at the RSVP module through local software processes and router configuration, other than RSVP neighbor adjacency. Junos OS provides security controls on access to the router and router configuration.

- Communication with external RSVP neighbors

RSVP signaled MPLS LSPs depend on the services of RSVP and IGP to communicate RSVP messages among neighboring routers across the network. Because the RSVP sessions involve communication outside of the local router, they are subject to many forms of attack, such as spoofing of peers, injection of falsified RSVP messages and route updates, and attacks on the underlying TCP/UDP transport for sessions. Junos OS provides countermeasures for such attack vectors.

- Resource limits and denial of service

Junos OS provides several mechanisms through policers and filters to protect against denial-of-service attacks based on injecting higher than the expected traffic demands. At the MPLS LSP level, Junos OS allows operators to configure limits on the LSP

bandwidth and the number of LSPs. However, like point-to-point RSVP LSPs, container LSPs do not enforce limits on the volume of traffic forwarded over these LSPs.

Supported and Unsupported Features

Junos OS supports the following container LSP features:

- Equal-bandwidth-based LSP splitting mechanism
- Aggregate-bandwidth-based LSP splitting and merging in a make-before-break way
- LSP-number-based naming mechanism for dynamically created member LSPs
- Periodic sampling mechanisms to estimate aggregate bandwidth
- Interoperability with auto-bandwidth feature
- ECMP using the dynamically created LSPs
- LDP-tunneling on the dynamically created LSP
- Configuring container LSP using IGP shortcuts
- Aggregated Ethernet links
- Logical systems

Junos OS does **not** support the following container LSP functionality:

- Node and link disjoint paths for different LSPs between an ingress and an egress routing device
- Bandwidth allocation policy different from equal bandwidth policy at the normalization event
- Constraint-based routing path computation to find equal IGP cost paths for different LSPs
- RSVP objects, such as **MLSP_TUNNEL Sender Template**, and **MLSP_TUNNEL Filter Specification** defined in [KOMPELLA-MLSP]
- Change in topology as a trigger for LSP splitting and merging
- Change in topology and link failure as a trigger for normalization, unless member LSPs go down
- Egress protection on container LSP
- Container LSP as a backup LSP for IGP interface
- Container LSP configured as IGP interface as forwarding address
- Container LSP as provider tunnel for multicast VPNs
- Dynamic LSPs for normalization
- CCC using container LSP
- Secondary paths for container LSP
- Bidirectional container LSP

- Policing
- Static routes using container LSPs as next hops on a best-effort basis
- External path computing entity, such as PCE
- Graceful Routing Engine switchover
- Nonstop active routing
- Unified ISSU
- Multichassis
- IPv6

Related Documentation

- [Example: Configuring Dynamic Bandwidth Management Using Container LSP on page 413](#)
- [Maximize Bandwidth Utilization with Juniper Networks TE++](#)

Configuring Dynamic Bandwidth Management Using Container LSP

You can configure a container LSP to enable load balancing across multiple point-to-point LSPs dynamically. A container LSP includes one or more member LSPs between the same ingress and egress routing devices. The member LSPs are similar to independent point-to-point LSPs, and each member LSP takes a different path to the same destination and can be routed along a different IGP cost path.

A container LSP provides support for dynamic bandwidth management by enabling the ingress router to dynamically add and remove member LSPs through a process called LSP splitting and LSP merging, respectively, based on configuration and aggregate traffic. Besides addition and deletion, member LSPs can also be re-optimized with different bandwidth values in a make-before-break way.

Before you begin:

1. Configure the device interfaces.
2. Configure the device router ID and autonomous system number.
3. Configure the following protocols:
 - RSVP
 - BGP
Configure a BGP group to peer device with remote provider edge (PE) device.
 - OSPF
Enable traffic engineering capabilities.
4. Configure a VRF routing instance.

To configure the PE device:

1. Enable MPLS on all the interfaces (excluding the management interface).

```
[edit protocols]
user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable
```

2. Configure the MPLS statistics parameters.

```
[edit protocols]
user@PE1# set mpls statistics file file-name
user@PE1# set mpls statistics file size size
user@PE1# set mpls statistics interval seconds
user@PE1# set mpls statistics auto-bandwidth
```

3. Configure the label-switched path (LSP) template parameters.

```
[edit protocols]
user@PE1# set mpls label-switched-path template-name template
user@PE1# set mpls label-switched-path template-name optimize-timer seconds
user@PE1# set mpls label-switched-path template-name link-protection
user@PE1# set mpls label-switched-path template-name adaptive
user@PE1# set mpls label-switched-path template-name auto-bandwidth
    adjust-interval seconds
user@PE1# set mpls label-switched-path template-name auto-bandwidth
    adjust-threshold seconds
user@PE1# set mpls label-switched-path template-name auto-bandwidth
    minimum-bandwidth mbps
user@PE1# set mpls label-switched-path template-name auto-bandwidth
    maximum-bandwidth mbps
```

4. Configure a container LSP between the two PE routers, and assign the LSP template.

```
[edit protocols]
user@PE1# set mpls container-label-switched-path container-lsp-name to
    remote-PE-ip-address
user@PE1# set mpls container-label-switched-path container-lsp-name
    label-switched-path-template template-name
```

5. Configure the container LSP parameters.

```
[edit protocols]
user@PE1# set mpls container-label-switched-path container-lsp-name
    splitting-merging maximum-member-lsps number
user@PE1# set mpls container-label-switched-path container-lsp-name
    splitting-merging minimum-member-lsps number
user@PE1# set mpls container-label-switched-path container-lsp-name
    splitting-merging splitting-bandwidth mbps
user@PE1# set mpls container-label-switched-path container-lsp-name
    splitting-merging merging-bandwidth mbps
user@PE1# set mpls container-label-switched-path container-lsp-name
    splitting-merging maximum-signaling-bandwidth mbps
user@PE1# set mpls container-label-switched-path container-lsp-name
    splitting-merging minimum-signaling-bandwidth mbps
user@PE1# set mpls container-label-switched-path container-lsp-name
    splitting-merging normalization normalize-interval seconds
user@PE1# set mpls container-label-switched-path container-lsp-name
    splitting-merging normalization failover-normalization
```



```

user@PE1# set mpls container-label-switched-path container-lsp-name
splitting-merging normalization normalization-retry-duration seconds
user@PE1# set mpls container-label-switched-path container-lsp-name
splitting-merging normalization normalization-retry-limits number
user@PE1# set mpls container-label-switched-path container-lsp-name
splitting-merging sampling cut-off-threshold number
user@PE1# set mpls container-label-switched-path container-lsp-name
splitting-merging sampling use-percentile number

```

6. Configure the policy statement to load-balance traffic.

```

[edit policy-options]
user@PE1# set policy-statement first-policy-name term 1 from protocol direct
user@PE1# set policy-statement first-policy-name term 1 then accept
user@PE1# set policy-statement second-policy-name then load-balance per-packet

```



NOTE: The policy to load-balance traffic should be assigned to the forwarding table configuration under the [edit routing-options] hierarchy level.

```

user@PE1# set forwarding-table export pplb

```

7. Verify and commit the configuration.

For example:

```

[edit protocols]
user@PE1# set rsvp preemption aggressive
user@PE1# set rsvp interface all aggregate
user@PE1# set rsvp interface fxp0.0 disable
user@PE1# set rsvp interface ge-0/0/1.0
user@PE1# set rsvp interface ge-0/0/2.0
user@PE1# set mpls statistics file auto-bw
user@PE1# set mpls statistics file size 10m
user@PE1# set mpls statistics interval 10
user@PE1# set mpls statistics auto-bandwidth
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 template
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 optimize-timer 30
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 link-protection
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 adaptive
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth
adjust-interval 300
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth
adjust-threshold 5
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth
minimum-bandwidth 10m
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth
maximum-bandwidth 10m
user@PE1# set mpls label-switched-path PE1-PE2-template-1 template
user@PE1# set mpls label-switched-path PE1-PE2-template-1 auto-bandwidth
adjust-interval 8000
user@PE1# set mpls label-switched-path PE1-PE2-template-1 auto-bandwidth
minimum-bandwidth 5m
user@PE1# set mpls label-switched-path PE1-PE2-template-1 auto-bandwidth
maximum-bandwidth 10m

```

```

user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
label-switched-path-template PE1-to-PE2-template1
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 to
10.255.102.128
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging maximum-member-lsps 20
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging minimum-member-lsps 2
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging splitting-bandwidth 40m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging merging-bandwidth 6m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging maximum-signaling-bandwidth 10m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging minimum-signaling-bandwidth 10m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging normalization normalize-interval 400
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging normalization failover-normalization
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging normalization normalization-retry-duration 20
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging normalization normalization-retry-limits 3
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging sampling cut-off-threshold 1
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging sampling use-percentile 90
user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable
user@PE1# set bgp group to-PE2 type internal
user@PE1# set bgp group to-PE2 local-address 10.255.102.166
user@PE1# set bgp group to-PE2 family inet-vpn unicast
user@PE1# set bgp group to-PE2 export direct
user@PE1# set bgp group to-PE2 neighbor 10.255.102.128
user@PE1# set ospf traffic-engineering
user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE1# set ospf area 0.0.0.0 interface ge-0/0/2.0 metric 100

[edit policy-options]
user@PE1# set policy-statement direct term 1 from protocol direct
user@PE1# set policy-statement direct term 1 then accept
user@PE1# set policy-statement pplb then load-balance per-packet

[edit]
user@PE1# commit
commit complete

```

- Related Documentation**
- [Dynamic Bandwidth Management Using Container LSP Overview on page 382](#)
 - [Example: Configuring Dynamic Bandwidth Management Using Container LSP on page 413](#)

Example: Configuring Dynamic Bandwidth Management Using Container LSP

This example shows how to enable dynamic bandwidth management by configuring container label-switched paths (LSPs) that enable load balancing across multiple point-to-point member LSPs.

- [Requirements on page 413](#)
- [Overview on page 413](#)
- [Configuration on page 414](#)
- [Verification on page 422](#)

Requirements

This example uses the following hardware and software components:

- Five routers that can be a combination of M Series, MX Series, or T Series routers, out of which two routers are provider edge (PE) routers and three routers are provider (P) routers
- Junos OS Release 14.2 or later running on all the routers

Before you begin:

1. Configure the device interfaces.
2. Configure the autonomous system numbers and router IDs for the devices.
3. Configure the following protocols:
 - RSVP
 - MPLS
 - BGP
 - OSPF

Overview

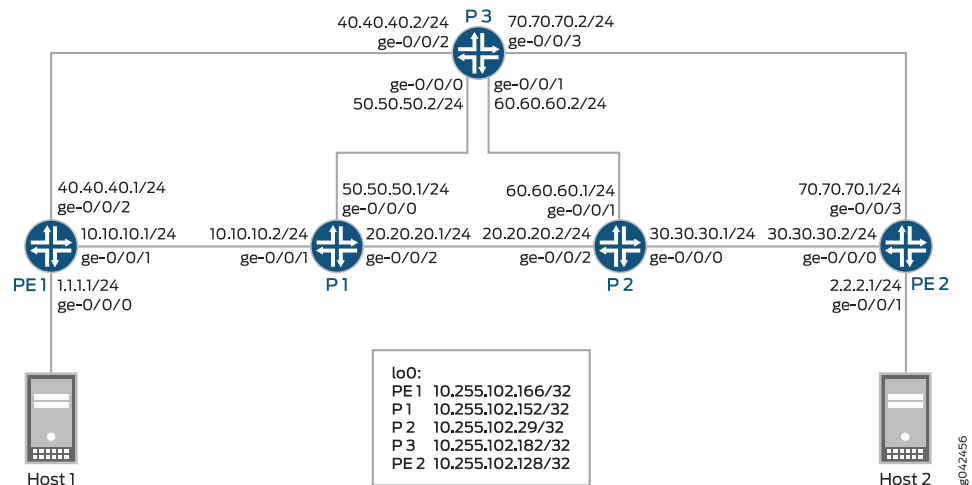
Starting with Junos OS Release 14.2, a new type of LSP, called a container LSP, is introduced to enable load balancing across multiple point-to-point LSPs. A container LSP includes one or more member LSPs between the same ingress and egress routing devices. The member LSPs are similar to independent point-to-point LSPs, and each member LSP takes a different path to the same destination and can be routed along a different IGP cost path.

A container LSP provides support for dynamic bandwidth management by enabling the ingress router to dynamically add and remove member LSPs through a process called LSP splitting and LSP merging, respectively, based on configuration and aggregate traffic. Besides addition and deletion, member LSPs can also be re-optimized with different bandwidth values in a make-before-break way.

Topology

Figure 37 on page 414 is a sample topology configured with container LSPs.

Figure 37: Dynamic Bandwidth Management Using Container LSP



In this example, Routers PE1 and PE2 are the PE routers connected to hosts Host1 and Host2, respectively. The core routers, Routers P1, P2, and P3 connect to the PE routers.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```

PE1 set interfaces ge-0/0/0 unit 0 family inet address 1.1.1/24
    set interfaces ge-0/0/1 unit 0 family inet address 10.10.10.1/24
    set interfaces ge-0/0/1 unit 0 family mpls
    set interfaces ge-0/0/2 unit 0 family inet address 40.40.40.1/24
    set interfaces ge-0/0/2 unit 0 family mpls
    set interfaces lo0 unit 0 family inet address 10.255.102.166/32
    set interfaces lo0 unit 0 family mpls
    set routing-options router-id 10.255.102.166
    set routing-options autonomous-system 1234
    set routing-options forwarding-table export pplib
    set protocols rsvp preemption aggressive
    set protocols rsvp interface all aggregate
    set protocols rsvp interface fxp0.0 disable
    set protocols rsvp interface ge-0/0/1.0
    set protocols rsvp interface ge-0/0/2.0
    set protocols mpls statistics file auto-bw
    set protocols mpls statistics file size 10m
    set protocols mpls statistics interval 10
    set protocols mpls statistics auto-bandwidth
    set protocols mpls label-switched-path PE1-to-PE2-template1 template
  
```

```
set protocols mpls label-switched-path PE1-to-PE2-template1 optimize-timer 30
set protocols mpls label-switched-path PE1-to-PE2-template1 link-protection
set protocols mpls label-switched-path PE1-to-PE2-template1 adaptive
set protocols mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth
  adjust-interval 300
set protocols mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth
  adjust-threshold 5
set protocols mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth
  minimum-bandwidth 10m
set protocols mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth
  maximum-bandwidth 10m
set protocols mpls container-label-switched-path PE1-PE2-container-100
  label-switched-path-template PE1-to-PE2-template1
set protocols mpls container-label-switched-path PE1-PE2-container-100 to
  10.255.102.128
set protocols mpls container-label-switched-path PE1-PE2-container-100
  splitting-merging maximum-member-lsps 20
set protocols mpls container-label-switched-path PE1-PE2-container-100
  splitting-merging minimum-member-lsps 2
set protocols mpls container-label-switched-path PE1-PE2-container-100
  splitting-merging splitting-bandwidth 40m
set protocols mpls container-label-switched-path PE1-PE2-container-100
  splitting-merging merging-bandwidth 6m
set protocols mpls container-label-switched-path PE1-PE2-container-100
  splitting-merging maximum-signaling-bandwidth 10m
set protocols mpls container-label-switched-path PE1-PE2-container-100
  splitting-merging minimum-signaling-bandwidth 10m
set protocols mpls container-label-switched-path PE1-PE2-container-100
  splitting-merging normalization normalize-interval 400
set protocols mpls container-label-switched-path PE1-PE2-container-100
  splitting-merging normalization failover-normalization
set protocols mpls container-label-switched-path PE1-PE2-container-100
  splitting-merging normalization normalization-retry-duration 20
set protocols mpls container-label-switched-path PE1-PE2-container-100
  splitting-merging normalization normalization-retry-limits 3
set protocols mpls container-label-switched-path PE1-PE2-container-100
  splitting-merging sampling cut-off-threshold 1
set protocols mpls container-label-switched-path PE1-PE2-container-100
  splitting-merging sampling use-percentile 90
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group to-PE2 type internal
set protocols bgp group to-PE2 local-address 10.255.102.166
set protocols bgp group to-PE2 family inet-vpn unicast
set protocols bgp group to-PE2 export direct
set protocols bgp group to-PE2 neighbor 10.255.102.128
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface ge-0/0/2.0 metric 100
set policy-options policy-statement direct term 1 from protocol direct
set policy-options policy-statement direct term 1 then accept
set policy-options policy-statement pplb then load-balance per-packet
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-0/0/0.0
set routing-instances vpn1 route-distinguisher 10.255.102.166:1
```

```
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label

P1  set interfaces ge-0/0/0 unit 0 family inet address 50.50.50.1/24
    set interfaces ge-0/0/0 unit 0 family mpls
    set interfaces ge-0/0/1 unit 0 family inet address 10.10.10.2/24
    set interfaces ge-0/0/1 unit 0 family mpls
    set interfaces ge-0/0/2 unit 0 family inet address 20.20.20.1/24
    set interfaces ge-0/0/2 unit 0 family mpls
    set interfaces lo0 unit 0 family inet address 10.255.102.152/32
    set protocols rsvp interface all aggregate
    set protocols rsvp interface fxp0.0 disable
    set protocols mpls interface all
    set protocols mpls interface fxp0.0 disable
    set protocols ospf traffic-engineering
    set protocols ospf area 0.0.0.0 interface all
    set protocols ospf area 0.0.0.0 interface fxp0.0 disable
    set protocols ospf area 0.0.0.0 interface ge-0/0/0.0 metric 100

P2  set interfaces ge-0/0/0 unit 0 family inet address 30.30.30.1/24
    set interfaces ge-0/0/0 unit 0 family mpls
    set interfaces ge-0/0/1 unit 0 family inet address 60.60.60.1/24
    set interfaces ge-0/0/1 unit 0 family mpls
    set interfaces ge-0/0/2 unit 0 family inet address 20.20.20.2/24
    set interfaces ge-0/0/2 unit 0 family mpls
    set interfaces lo0 unit 0 family inet address 10.255.102.29/32
    set protocols rsvp interface all aggregate
    set protocols rsvp interface fxp0.0 disable
    set protocols mpls statistics file auto_bw
    set protocols mpls statistics file size 10m
    set protocols mpls statistics interval 5
    set protocols mpls statistics auto-bandwidth
    set protocols mpls icmp-tunneling
    set protocols mpls interface all
    set protocols mpls interface fxp0.0 disable
    set protocols ospf traffic-engineering
    set protocols ospf area 0.0.0.0 interface all
    set protocols ospf area 0.0.0.0 interface fxp0.0 disable
    set protocols ospf area 0.0.0.0 interface ge-0/0/1.0 metric 100

P3  set interfaces ge-0/0/0 unit 0 family inet address 50.50.50.2/24
    set interfaces ge-0/0/0 unit 0 family mpls
    set interfaces ge-0/0/1 unit 0 family inet address 60.60.60.2/24
    set interfaces ge-0/0/1 unit 0 family mpls
    set interfaces ge-0/0/2 unit 0 family inet address 40.40.40.2/24
    set interfaces ge-0/0/2 unit 0 family mpls
    set interfaces ge-0/0/3 unit 0 family inet address 70.70.70.2/24
    set interfaces ge-0/0/3 unit 0 family mpls
    set interfaces lo0 unit 0 family inet address 10.255.102.182/32
    set protocols rsvp interface all aggregate
    set protocols rsvp interface fxp0.0 disable
    set protocols mpls icmp-tunneling
    set protocols mpls interface all
    set protocols mpls interface fxp0.0 disable
    set protocols ospf traffic-engineering
```

```
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
```

```
PE2 set interfaces ge-0/0/0 unit 0 family inet address 30.30.30.2/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.1/24
set interfaces ge-0/0/3 unit 0 family inet address 70.70.70.1/24
set interfaces ge-0/0/3 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.102.128/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 10.255.102.128
set routing-options autonomous-system 1234
set protocols rsvp interface all aggregate
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group to-PE1 type internal
set protocols bgp group to-PE1 local-address 10.255.102.128
set protocols bgp group to-PE1 family inet-vpn unicast
set protocols bgp group to-PE1 neighbor 10.255.102.166
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set policy-options policy-statement direct from protocol direct
set policy-options policy-statement direct then accept
set routing-instances vpn1 instance-type vrf
set routing-instances vpn1 interface ge-0/0/1.0
set routing-instances vpn1 route-distinguisher 10.255.102.128:1
set routing-instances vpn1 vrf-target target:1:1
set routing-instances vpn1 vrf-table-label
```

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router PE1:

1. Configure the Router PE1 interfaces.

```
[edit interfaces]
user@PE1# set ge-0/0/0 unit 0 family inet address 1.1.1.1/24

user@PE1# set ge-0/0/1 unit 0 family inet address 10.10.10.1/24
user@PE1# set ge-0/0/1 unit 0 family mpls

user@PE1# set ge-0/0/2 unit 0 family inet address 40.40.40.1/24
user@PE1# set ge-0/0/2 unit 0 family mpls

user@PE1# set lo0 unit 0 family inet address 10.255.102.166/32
user@PE1# set lo0 unit 0 family mpls
```

2. Configure the router ID and autonomous system number for Router PE1.

```
[edit routing-options]
user@PE1# set router-id 10.255.102.166
```

- ```
user@PE1# set autonomous-system 1234
```
3. Enable the policy to load-balance traffic.  

```
[edit routing-options]
user@PE1# set forwarding-table export pplb
```
  4. Enable RSVP on all Router PE1 interfaces (excluding the management interface).  

```
[edit protocols]
user@PE1# set rsvp preemption aggressive
user@PE1# set rsvp interface all aggregate
user@PE1# set rsvp interface fxp0.0 disable
user@PE1# set rsvp interface ge-0/0/1.0
user@PE1# set rsvp interface ge-0/0/2.0
```
  5. Enable MPLS on all the interfaces of Router PE1 (excluding the management interface).  

```
[edit protocols]
user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable
```
  6. Configure the MPLS statistics parameters.  

```
[edit protocols]
user@PE1# set mpls statistics file auto-bw
user@PE1# set mpls statistics file size 10m
user@PE1# set mpls statistics interval 10
user@PE1# set mpls statistics auto-bandwidth
```
  7. Configure the label-switched path (LSP) template parameters.  

```
[edit protocols]
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 template
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 optimize-timer 30
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 link-protection
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 adaptive
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth
adjust-interval 300
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth
adjust-threshold 5
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth
minimum-bandwidth 10m
user@PE1# set mpls label-switched-path PE1-to-PE2-template1 auto-bandwidth
maximum-bandwidth 10m
```
  8. Configure a container LSP between Router PE1 and Router PE2, and assign the PE1-to-PE2-template1 LSP template.  

```
[edit protocols]
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100 to
10.255.102.128
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
label-switched-path-template PE1-to-PE2-template1
```
  9. Configure the container LSP parameters.  

```
[edit protocols]
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging maximum-member-lsps 20
```



```

user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging minimum-member-lsps 2
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging splitting-bandwidth 40m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging merging-bandwidth 6m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging maximum-signaling-bandwidth 10m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging minimum-signaling-bandwidth 10m
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging normalization normalize-interval 400
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging normalization failover-normalization
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging normalization normalization-retry-duration 20
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging normalization normalization-retry-limits 3
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging sampling cut-off-threshold 1
user@PE1# set mpls container-label-switched-path PE1-PE2-container-100
splitting-merging sampling use-percentile 90

```

10. Configure the BGP group, and assign the local and neighbor IP addresses.

```

[edit protocols]
user@PE1# set bgp group to-PE2 type internal
user@PE1# set bgp group to-PE2 local-address 10.255.102.166
user@PE1# set bgp group to-PE2 neighbor 10.255.102.128
user@PE1# set bgp group to-PE2 family inet-vpn unicast
user@PE1# set bgp group to-PE2 export direct

```

11. Enable OSPF on all the interfaces of Router PE1 (excluding the management interface) along with traffic engineering capabilities.

```

[edit protocols]
user@PE1# set ospf traffic-engineering
user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE1# set ospf area 0.0.0.0 interface ge-0/0/2.0 metric 100

```

12. Configure the policy statement to load-balance traffic.

```

[edit policy-options]
user@PE1# set policy-statement direct term 1 from protocol direct
user@PE1# set policy-statement direct term 1 then accept
user@PE1# set policy-statement pplb then load-balance per-packet

```

13. Configure a routing instance on Router PE1, and assign the routing instance interface.

```

[edit routing-instances]
user@PE1# set vpn1 instance-type vrf
user@PE1# set vpn1 interface ge-0/0/0.0

```

14. Configure the route distinguisher, vrf target, and vrf-table label values for the VRF routing instance.

```

[edit routing-instances]
user@PE1# set vpn1 route-distinguisher 10.255.102.166:1

```

```
user@PE1# set vpn1 vrf-target target:1:1
user@PE1# set vpn1 vrf-table-label
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show interfaces
ge-0/0/0 {
 unit 0 {
 family inet {
 address 1.1.1.1/24;
 }
 }
}
ge-0/0/1 {
 unit 0 {
 family inet {
 address 10.10.10.1/24;
 }
 family mpls;
 }
}
ge-0/0/2 {
 unit 0 {
 family inet {
 address 40.40.40.1/24;
 }
 family mpls;
 }
}
lo0 {
 unit 0 {
 family inet {
 address 10.255.102.166/32;
 }
 family mpls;
 }
}

user@PE1# show routing-options
router-id 10.255.102.166;
autonomous-system 1234;
forwarding-table {
 export pplb;
}

user@PE1# show protocols
rsvp {
 preemption aggressive;
 interface all {
 aggregate;
 }
 interface fxp0.0 {
 disable;
 }
}
```

```

 }
 interface ge-0/0/1.0;
 interface ge-0/0/2.0;
 }
 mpls {
 statistics {
 file auto-bw size 10m;
 interval 10;
 auto-bandwidth;
 }
 label-switched-path PE1-to-PE2-template1 {
 template;
 optimize-timer 30;
 link-protection;
 adaptive;
 auto-bandwidth {
 adjust-interval 300;
 adjust-threshold 5;
 minimum-bandwidth 10m;
 maximum-bandwidth 10m;
 }
 }
 }
 container-label-switched-path PE1-PE2-container-100 {
 label-switched-path-template {
 PE1-to-PE2-template1;
 }
 to 10.255.102.128;
 splitting-merging {
 maximum-member-lsps 20;
 minimum-member-lsps 2;
 splitting-bandwidth 40m;
 merging-bandwidth 6m;
 maximum-signaling-bandwidth 10m;
 minimum-signaling-bandwidth 10m;
 normalization {
 normalize-interval 400;
 failover-normalization;
 normalization-retry-duration 20;
 normalization-retry-limits 3;
 }
 sampling {
 cut-off-threshold 1;
 use-percentile 90;
 }
 }
 }
 interface all;
 interface fxp0.0 {
 disable;
 }
}
bgp {
 group to-PE2 {
 type internal;
 local-address 10.255.102.166;
 family inet-vpn {

```

```
 unicast;
 }
 export direct;
 neighbor 10.255.102.128;
}
}
ospf {
 traffic-engineering;
 area 0.0.0.0 {
 interface all;
 interface fxp0.0 {
 disable;
 }
 interface ge-0/0/2.0 {
 metric 100;
 }
 }
}
}

user@PE1# show policy-options
policy-statement direct {
 term 1 {
 from protocol direct;
 then accept;
 }
}
policy-statement pplb {
 then load-balance {
 per-packet;
 }
}

user@PE1# show routing-instances
vpn1 {
 instance-type vrf;
 interface ge-0/0/0.0;
 route-distinguisher 10.255.102.166:1;
 vrf-target target:1:1;
 vrf-table-label;
}
```

## Verification

Confirm that the configuration is working properly.

- [Verifying the Container LSP Status Without Bandwidth on page 423](#)
- [Verifying the Container LSP Status with Increased Bandwidth \(Before Normalization\) on page 426](#)
- [Verifying the Container LSP Status with Increased Bandwidth \(After Normalization\) on page 428](#)
- [Verifying the Container LSP Splitting Process on page 432](#)
- [Verifying the Container LSP Statistics on page 432](#)

- [Verifying the Container LSP Status with Decreased Bandwidth \(Before Normalization\) on page 432](#)
- [Verifying the Container LSP Status with Decreased Bandwidth \(After Normalization\) on page 433](#)
- [Verifying the Container LSP Merging Process on page 434](#)
- [Verifying Failover Normalization on page 434](#)
- [Verifying Incremental Normalization on page 435](#)

---

### **[Verifying the Container LSP Status Without Bandwidth](#)**

---

**Purpose**    Verify the status of the container LSP.

**Action** From operational mode, run the **show mpls container-lsp extensive** command.

```

user@PE1> show mpls container-lsp extensive
Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 2
Normalization
 Min LSPs: 2, Max LSPs: 20
 Aggregate bandwidth: 20Mbps, Sampled Aggregate bandwidth: 0bps
 NormalizeTimer: 400 secs, NormalizeThreshold: 10%
 Max Signaling BW: 10Mbps, Min Signaling BW: 10Mbps, Splitting BW: 40Mbps, Merging
 BW: 6Mbps
 Mode: incremental-normalization, failover-normalization
 Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
 Normalization in 143 second(s)
 36 Jun 5 04:12:17.497 Clear history and statistics: on container
 (PE1-PE2-container-100)
 35 Jun 5 04:12:17.497 Avoid normalization: not needed with bandwidth 0 bps
 34 Jun 5 04:05:37.484 Clear history and statistics: on container
 (PE1-PE2-container-100)
 33 Jun 5 04:05:37.484 Avoid normalization: not needed with bandwidth 0 bps
 32 Jun 5 03:58:57.470 Clear history and statistics: on container
 (PE1-PE2-container-100)
 31 Jun 5 03:58:57.470 Avoid normalization: not needed with bandwidth 0 bps
 30 Jun 5 03:52:17.455 Clear history and statistics: on container
 (PE1-PE2-container-100)
 29 Jun 5 03:52:17.455 Avoid normalization: not needed with bandwidth 0 bps
 28 Jun 5 03:45:37.440 Clear history and statistics: on container
 (PE1-PE2-container-100)
 27 Jun 5 03:45:37.440 Avoid normalization: not needed with bandwidth 0 bps
 26 Jun 5 03:38:59.013 Normalization complete: container (PE1-PE2-container-100)
 with 2 members
 25 Jun 5 03:38:57.423 Delete member LSPs: PE1-PE2-container-100-3 through
 PE1-PE2-container-100-7
 24 Jun 5 03:38:57.423 Normalize: container (PE1-PE2-container-100) create 2
 LSPs, min bw 10000000bps, member count 7
 23 Jun 5 03:38:57.423 Normalize: normalization with aggregate bandwidth 0 bps

 22 Jun 5 03:32:19.019 Normalization complete: container (PE1-PE2-container-100)
 with 7 members
 21 Jun 5 03:32:17.404 Clear history and statistics: on container
 (PE1-PE2-container-100)
 20 Jun 5 03:32:17.403 Normalize: container (PE1-PE2-container-100) into 7
 members - each with bandwidth 10000000 bps
 19 Jun 5 03:32:17.403 Normalize: normalization with aggregate bandwidth
 62914560 bps
 18 Jun 5 03:32:17.403 Normalize: normalizaton with 62914560 bps
 17 Jun 5 03:32:09.219 Normalization complete: container (PE1-PE2-container-100)
 with 7 members
 16 Jun 5 03:32:07.600 Clear history and statistics: on container
 (PE1-PE2-container-100)
 15 Jun 5 03:32:07.600 Normalize: container (PE1-PE2-container-100) into 7
 members - each with bandwidth 10000000 bps
 14 Jun 5 03:32:07.599 Normalize: normalization with aggregate bandwidth
 62914560 bps
 13 Jun 5 03:32:07.599 Normalize: normalizaton with 62914560 bps
 12 Jun 5 03:26:57.295 Clear history and statistics: on container
 (PE1-PE2-container-100)
 11 Jun 5 03:26:57.295 Avoid normalization: not needed with bandwidth 0 bps
 10 Jun 5 03:20:18.297 Normalization complete: container (PE1-PE2-container-100)
 with 2 members

```

```

 9 Jun 5 03:20:17.281 Normalize: container (PE1-PE2-container-100) create 2
LSPs, min bw 10000000bps, member count 0
 8 Jun 5 03:20:17.281 Normalize: normalization with aggregate bandwidth 0 bps

 7 Jun 5 03:17:43.218 Clear history and statistics: on container
(PE1-PE2-container-100)
 6 Jun 5 03:17:43.218 Avoid normalization: not needed with bandwidth 0 bps
 5 Jun 5 03:17:43.212 Normalize: container (PE1-PE2-container-100) received
PathErr on member PE1-PE2-container-100-2
 4 Jun 5 03:17:43.212 Normalize: container (PE1-PE2-container-100) received
PathErr on member PE1-PE2-container-100-1
 3 Jun 5 03:12:47.323 Normalization complete: container (PE1-PE2-container-100)
with 2 members
 2 Jun 5 03:12:16.555 Normalize: container (PE1-PE2-container-100) create 2
LSPs, min bw 10000000bps, member count 0
 1 Jun 5 03:12:16.555 Normalize: normalization with aggregate bandwidth 0 bps

10.255.102.128
 From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-1

 ActivePath: (primary)
 LSPTYPE: Dynamic Configured, Penultimate hop popping
 LoadBalance: Random
 Autobandwidth
 MinBW: 10Mbps, MaxBW: 10Mbps
 AdjustTimer: 300 secs
 Max AvgBW util: 0bps, Bandwidth Adjustment in 12 second(s).
 Overflow limit: 0, Overflow sample count: 0
 Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary State: Up
 Priorities: 7 0
 Bandwidth: 10Mbps
 SmartOptimizeTimer: 180
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
 10.10.10.2 20.20.20.2 30.30.30.2
 17 Jun 5 03:38:59.013 Make-before-break: Switched to new instance
 16 Jun 5 03:38:58.003 Record Route: 10.10.10.2 20.20.20.2 30.30.30.2
 15 Jun 5 03:38:58.003 Up
 14 Jun 5 03:38:57.423 Originate make-before-break call
 13 Jun 5 03:38:57.423 CSPF: computation result accepted 10.10.10.2 20.20.20.2
30.30.30.2
 12 Jun 5 03:33:30.400 CSPF: computation result ignored, new path no benefit
 11 Jun 5 03:32:17.403 Pending old path instance deletion
 10 Jun 5 03:32:09.218 Make-before-break: Switched to new instance
 9 Jun 5 03:32:08.202 Record Route: 10.10.10.2 20.20.20.2 30.30.30.2
 8 Jun 5 03:32:08.202 Up
 7 Jun 5 03:32:07.603 Originate make-before-break call
 6 Jun 5 03:32:07.603 CSPF: computation result accepted 10.10.10.2 20.20.20.2
30.30.30.2
 5 Jun 5 03:20:18.278 Selected as active path
 4 Jun 5 03:20:18.277 Record Route: 10.10.10.2 20.20.20.2 30.30.30.2
 3 Jun 5 03:20:18.277 Up
 2 Jun 5 03:20:17.281 Originate Call
 1 Jun 5 03:20:17.281 CSPF: computation result accepted 10.10.10.2 20.20.20.2
30.30.30.2
 Created: Thu Jun 5 03:20:16 2014

```

```

10.255.102.128
 From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPName: PE1-PE2-container-100-2

 ActivePath: (primary)
 LSPType: Dynamic Configured, Penultimate hop popping
 LoadBalance: Random
 Autobandwidth
 MinBW: 10Mbps, MaxBW: 10Mbps
 AdjustTimer: 300 secs
 Max AvgBW util: 0bps, Bandwidth Adjustment in 76 second(s).
 Overflow limit: 0, Overflow sample count: 0
 Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary State: Up
 Priorities: 7 0
 Bandwidth: 10Mbps
 SmartOptimizeTimer: 180
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
 10.10.10.2 20.20.20.2 30.30.30.2
 17 Jun 5 03:38:59.013 Make-before-break: Switched to new instance
 16 Jun 5 03:38:58.002 Record Route: 10.10.10.2 20.20.20.2 30.30.30.2
 15 Jun 5 03:38:58.002 Up
 14 Jun 5 03:38:57.423 Originate make-before-break call
 13 Jun 5 03:38:57.423 CSPF: computation result accepted 10.10.10.2 20.20.20.2
30.30.30.2
 12 Jun 5 03:33:26.189 CSPF: computation result ignored, new path no benefit
 11 Jun 5 03:32:17.403 Pending old path instance deletion
 10 Jun 5 03:32:09.219 Make-before-break: Switched to new instance
 9 Jun 5 03:32:08.204 Record Route: 10.10.10.2 20.20.20.2 30.30.30.2
 8 Jun 5 03:32:08.204 Up
 7 Jun 5 03:32:07.603 Originate make-before-break call
 6 Jun 5 03:32:07.603 CSPF: computation result accepted 10.10.10.2 20.20.20.2
30.30.30.2
 5 Jun 5 03:20:18.297 Selected as active path
 4 Jun 5 03:20:18.295 Record Route: 10.10.10.2 20.20.20.2 30.30.30.2
 3 Jun 5 03:20:18.295 Up
 2 Jun 5 03:20:17.281 Originate Call
 1 Jun 5 03:20:17.281 CSPF: computation result accepted 10.10.10.2 20.20.20.2
30.30.30.2
 Created: Thu Jun 5 03:20:16 2014
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

**Meaning** The container LSP is established between Routers PE1 and PE2.

### Verifying the Container LSP Status with Increased Bandwidth (Before Normalization)

**Purpose** Verify the status of the container LSP with increased bandwidth before normalization happens.



**Action** From operational mode, run the **show mpls container-lsp extensive** command.

```

user@PE1> show mpls container-lsp extensive
Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 2
Normalization
 Min LSPs: 2, Max LSPs: 20
 Aggregate bandwidth: 20Mbps, Sampled Aggregate bandwidth: 42.6984Mbps
 NormalizeTimer: 400 secs, NormalizeThreshold: 10%
 Max Signaling BW: 10Mbps, Min Signaling BW: 10Mbps, Splitting BW: 40Mbps, Merging
 BW: 6Mbps
 Mode: incremental-normalization, failover-normalization
 Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
 Normalization in 321 second(s)
 3 Jun 5 21:22:34.731 Normalization complete: container (PE1-PE2-container-100)
 with 2 members
 2 Jun 5 21:22:15.503 Normalize: container (PE1-PE2-container-100) create 2
 LSPs, min bw 100000000bps, member count 0
 1 Jun 5 21:22:15.503 Normalize: normalization with aggregate bandwidth 0 bps

10.255.102.128
 From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-1

 ActivePath: (primary)
 Link protection desired
 LSPTYPE: Dynamic Configured, Penultimate hop popping
 LoadBalance: Random
 Autobandwidth
 MinBW: 10Mbps, MaxBW: 10Mbps
 AdjustTimer: 300 secs AdjustThreshold: 5%
 Max AvgBW util: 23.9893Mbps, Bandwidth Adjustment in 221 second(s).
 Overflow limit: 0, Overflow sample count: 6
 Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary State: Up
 Priorities: 7 0
 Bandwidth: 10Mbps
 OptimizeTimer: 30
 SmartOptimizeTimer: 180
 Reoptimization in 9 second(s).
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
 10.255.102.166(flag=0x20) 10.10.10.2(Label=303440)
 10.255.102.29(flag=0x20) 20.20.20.2(Label=302144) 10.255.102.128(flag=0x20)
 30.30.30.2(Label=3)

10.255.102.128
 From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-2

 ActivePath: (primary)
 Link protection desired
 LSPTYPE: Dynamic Configured, Penultimate hop popping
 LoadBalance: Random
 Autobandwidth
 MinBW: 10Mbps, MaxBW: 10Mbps
 AdjustTimer: 300 secs AdjustThreshold: 5%
 Max AvgBW util: 22.1438Mbps, Bandwidth Adjustment in 221 second(s).
 Overflow limit: 0, Overflow sample count: 6

```

```
Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
 Priorities: 7 0
 Bandwidth: 10Mbps
 OptimizeTimer: 30
 SmartOptimizeTimer: 180
 Reoptimization in 9 second(s).
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
 10.255.102.166(flag=0x20) 10.10.10.2(Label=303456)
10.255.102.29(flag=0x20) 20.20.20.2(Label=302160) 10.255.102.128(flag=0x20)
30.30.30.2(Label=3)

Total 2 displayed, Up 2, Down 0
```

**Meaning** Because normalization has not happened, the member LSP count remains at 2.

#### Verifying the Container LSP Status with Increased Bandwidth (After Normalization)

---

**Purpose** Verify the status of the container LSP with increased bandwidth after normalization happens.

**Action** From operational mode, run the **show mpls container-lsp extensive** command.

```

user@PE1> show mpls container-lsp extensive
Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 5
Normalization
 Min LSPs: 2, Max LSPs: 20
 Aggregate bandwidth: 50Mbps, Sampled Aggregate bandwidth: 45.8873Mbps
 NormalizeTimer: 400 secs, NormalizeThreshold: 10%
 Max Signaling BW: 10Mbps, Min Signaling BW: 10Mbps, Splitting BW: 40Mbps, Merging
 BW: 6Mbps
 Mode: incremental-normalization, failover-normalization
 Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
 Normalization in 169 second(s)
 7 Jun 5 21:29:02.921 Normalization complete: container (PE1-PE2-container-100)
 with 5 members
 6 Jun 5 21:28:55.505 Clear history and statistics: on container
 (PE1-PE2-container-100)
 5 Jun 5 21:28:55.505 Normalize: container (PE1-PE2-container-100) into 5
 members - each with bandwidth 10000000 bps
 4 Jun 5 21:28:55.504 Normalize: normalization with aggregate bandwidth
 45281580 bps
 3 Jun 5 21:22:34.731 Normalization complete: container (PE1-PE2-container-100)
 with 2 members
 2 Jun 5 21:22:15.503 Normalize: container (PE1-PE2-container-100) create 2
 LSPs, min bw 10000000bps, member count 0
 1 Jun 5 21:22:15.503 Normalize: normalization with aggregate bandwidth 0 bps

10.255.102.128
 From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-1

 ActivePath: (primary)
 Link protection desired
 LSPtype: Dynamic Configured, Penultimate hop popping
 LoadBalance: Random
 Autobandwidth
 MinBW: 10Mbps, MaxBW: 10Mbps
 AdjustTimer: 300 secs AdjustThreshold: 5%
 Max AvgBW util: 11.0724Mbps, Bandwidth Adjustment in 129 second(s).
 Overflow limit: 0, Overflow sample count: 1
 Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary State: Up
 Priorities: 7 0
 Bandwidth: 10Mbps
 OptimizeTimer: 30
 SmartOptimizeTimer: 180
 Reoptimization in 12 second(s).
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
 10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
 10.255.102.166(flag=0x20) 10.10.10.2(Label=303488)
 10.255.102.29(flag=0x20) 20.20.20.2(Label=302224) 10.255.102.128(flag=0x20)
 30.30.30.2(Label=3)

10.255.102.128
 From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-2

 ActivePath: (primary)

```

```

Link protection desired
LSPTYPE: Dynamic Configured, Penultimate hop popping
LoadBalance: Random
Autobandwidth
MinBW: 10Mbps, MaxBW: 10Mbps
AdjustTimer: 300 secs AdjustThreshold: 5%
Max AvgBW util: 8.50751Mbps, Bandwidth Adjustment in 189 second(s).
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 11, Underflow Max AvgBW: 8.50751Mbps

Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
 Priorities: 7 0
 Bandwidth: 10Mbps
 OptimizeTimer: 30
 SmartOptimizeTimer: 180
 Reoptimization in 6 second(s).
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
 10.255.102.166(flag=0x20) 10.10.10.2(Label=303504)
10.255.102.29(flag=0x20) 20.20.20.2(Label=302240) 10.255.102.128(flag=0x20)
30.30.30.2(Label=3)

10.255.102.128
From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPName: PE1-PE2-container-100-3

ActivePath: (primary)
Link protection desired
LSPTYPE: Dynamic Configured, Penultimate hop popping
LoadBalance: Random
Autobandwidth
MinBW: 10Mbps, MaxBW: 10Mbps
AdjustTimer: 300 secs AdjustThreshold: 5%
Max AvgBW util: 9.59422Mbps, Bandwidth Adjustment in 249 second(s).
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 5, Underflow Max AvgBW: 9.59422Mbps

Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
 Priorities: 7 0
 Bandwidth: 10Mbps
 OptimizeTimer: 30
 SmartOptimizeTimer: 180
 Reoptimization in 25 second(s).
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
 10.255.102.166(flag=0x20) 10.10.10.2(Label=303472)
10.255.102.29(flag=0x20) 20.20.20.2(Label=302176) 10.255.102.128(flag=0x20)
30.30.30.2(Label=3)

10.255.102.128
From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPName: PE1-PE2-container-100-4

ActivePath: (primary)
Link protection desired
LSPTYPE: Dynamic Configured, Penultimate hop popping
LoadBalance: Random

```

```

Autobandwidth
MinBW: 10Mbps, MaxBW: 10Mbps
AdjustTimer: 300 secs AdjustThreshold: 5%
Max AvgBW util: 9.16169Mbps, Bandwidth Adjustment in 9 second(s).
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 29, Underflow Max AvgBW: 9.16169Mbps

Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
 Priorities: 7 0
 Bandwidth: 10Mbps
 OptimizeTimer: 30
 SmartOptimizeTimer: 180
 Reoptimization in 1 second(s).
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
 10.255.102.166(flag=0x20) 10.10.10.2(Label=303520)
10.255.102.29(flag=0x20) 20.20.20.2(Label=302192) 10.255.102.128(flag=0x20)
30.30.30.2(Label=3)

10.255.102.128
 From: 10.255.102.166, State: Up, ActiveRoute: 0, LSPname: PE1-PE2-container-100-5

ActivePath: (primary)
Link protection desired
LSPTYPE: Dynamic Configured, Penultimate hop popping
LoadBalance: Random
Autobandwidth
MinBW: 10Mbps, MaxBW: 10Mbps
AdjustTimer: 300 secs AdjustThreshold: 5%
Max AvgBW util: 8.39908Mbps, Bandwidth Adjustment in 69 second(s).
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 23, Underflow Max AvgBW: 8.39908Mbps

Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
 Priorities: 7 0
 Bandwidth: 10Mbps
 OptimizeTimer: 30
 SmartOptimizeTimer: 180
 Reoptimization in 17 second(s).
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.10.10.2 S 20.20.20.2 S 30.30.30.2 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
 10.255.102.166(flag=0x20) 10.10.10.2(Label=303536)
10.255.102.29(flag=0x20) 20.20.20.2(Label=302208) 10.255.102.128(flag=0x20)
30.30.30.2(Label=3)
Total 5 displayed, Up 5, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

**Meaning** At the expiry of the normalization timer, the container LSP is split into five member LSPs, each with 10 Mbps (minimum and maximum signaling bandwidth). As a result, the

aggregate bandwidth is 50 Mbps.

### Verifying the Container LSP Splitting Process

**Purpose** Verify the container LSP splitting process after normalization happens.

**Action** From operational mode, run the **show route 2.2.2** command.

```
user@PE1> show route 2.2.2
vpn1.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.0/24 *[BGP/170] 00:12:14, localpref 100, from 10.255.102.128
 AS path: I, validation-state: unverified
>to 10.10.10.2 via ge-0/0/1.0, label-switched-path PE1-PE2-container100-1
to 10.10.10.2 via ge-0/0/1.0, label-switched-path PE1-PE2-container100-2
to 10.10.10.2 via ge-0/0/1.0, label-switched-path PE1-PE2-container100-3
to 10.10.10.2 via ge-0/0/1.0, label-switched-path PE1-PE2-container100-4
to 10.10.10.2 via ge-0/0/1.0, label-switched-path PE1-PE2-container100-5
```

**Meaning** After LSP splitting, Router PE1 has injected the forwarding adjacency.

### Verifying the Container LSP Statistics

**Purpose** Verify the container LSP statistics after normalization happens.

**Action** From operational mode, run the **show mpls container-lsp statistics** command.

```
user@PE1> show mpls container-lsp statistics
Ingress LSP: 1 sessions
Container LSP name State Member LSP count
PE1-PE2-container-100 Up 5
To From State Packets Bytes LSPname
10.255.102.128 10.255.102.166 Up 15166271 2062612856
PE1-PE2-container-100-1
10.255.102.128 10.255.102.166 Up 12289912 1671428032
PE1-PE2-container-100-2
10.255.102.128 10.255.102.166 Up 13866911 1885899896
PE1-PE2-container-100-3
10.255.102.128 10.255.102.166 Up 12558707 1707984152
PE1-PE2-container-100-4
10.255.102.128 10.255.102.166 Up 11512151 1565652536
PE1-PE2-container-100-5
```

**Meaning** Traffic is load-balanced across the newly created member LSPs.

### Verifying the Container LSP Status with Decreased Bandwidth (Before Normalization)

**Purpose** Verify the status of the container LSP with decreased bandwidth before normalization happens.

**Action** From operational mode, run the **show mpls container-lsp detail** command.

```
user@PE1> show mpls container-lsp detail
Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 5
Normalization
 Min LSPs: 2, Max LSPs: 20
 Aggregate bandwidth: 50Mbps, Sampled Aggregate bandwidth: 2.0215Mbps
 NormalizeTimer: 400 secs, NormalizeThreshold: 10%
 Max Signaling BW: 10Mbps, Min Signaling BW: 10Mbps, Splitting BW: 40Mbps, Merging
 BW: 6Mbps
 Mode: incremental-normalization, failover-normalization
 Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
 Normalization in 384 second(s)
---Output truncated---
```

**Meaning** Because normalization has not happened, the member LSP count remains at 5.

### Verifying the Container LSP Status with Decreased Bandwidth (After Normalization)

**Purpose** Verify the status of the container LSP with decreased bandwidth after normalization happens.

**Action** From operational mode, run the **show mpls container-lsp detail** command.

```
user@PE1> show mpls container-lsp detail
Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 2
Normalization
 Min LSPs: 2, Max LSPs: 20
 Aggregate bandwidth: 20Mbps, Sampled Aggregate bandwidth: 0bps
 NormalizeTimer: 400 secs, NormalizeThreshold: 10%
 Max Signaling BW: 10Mbps, Min Signaling BW: 10Mbps, Splitting BW: 40Mbps, Merging
 BW: 6Mbps
 Mode: incremental-normalization, failover-normalization
 Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
 Normalization in 397 second(s)
 22 Jun 5 22:30:37.094 Clear history and statistics: on container
 (PE1-PE2-container-100)
 21 Jun 5 22:30:37.094 Delete member LSPs: PE1-PE2-container-100-3 through
 PE1-PE2-container-100-5
 20 Jun 5 22:30:37.090 Normalize: container (PE1-PE2-container-100) into 2
 members - each with bandwidth 10000000 bps
 19 Jun 5 22:30:37.090 Normalize: normalization with aggregate bandwidth 2037595
 bps
 18 Jun 5 22:30:37.090 Normalize: normalization with 2037595 bps
---Output truncated---
```

**Meaning** At the expiry of the normalization timer, the container LSP merging takes place because there is an overall reduction in bandwidth. The member LSPs are merged, and the member LSP count is 2 after normalization.

### Verifying the Container LSP Merging Process

---

**Purpose** Verify the container LSP splitting process after normalization happens.

**Action** From operational mode, run the **show route 2.2.2** command.

```
user@PE1> show route 2.2.2
vpn1.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.0/24 *[BGP/170] 01:09:45, localpref 100, from 10.255.102.128
 AS path: I, validation-state: unverified
 > to 10.10.10.2 via ge-0/0/1.0, label-switched-path
PE1-PE2-container-100-1
 to 10.10.10.2 via ge-0/0/1.0, label-switched-path
PE1-PE2-container-100-2
```

**Meaning** After LSP merging, Router PE1 has deleted the merged member LSPs.

### Verifying Failover Normalization

---

**Purpose** Verify load redistribution when traffic is sent at 35 Mbps and the link between Routers P1 and P2 is disabled. Arrival of PathErr on link failure triggers immediate normalization.

To enable failover normalization, include the **failover-normalization** configuration statement at the **[edit protocols mpls container-label-switched-path container-lsp-name splitting-merging normalization]** hierarchy level.



**Action** From operational mode, run the **show mpls container-lsp** command.

```
user@PE1> show mpls container-lsp
Ingress LSP: 1 sessions
Container LSP name
PE1-PE2-container-100
To From State Rt P State Member LSP count
Up ActivePath 2
10.255.102.128 10.255.102.166 Up 0 *
PE1-PE2-container-100-1
10.255.102.128 10.255.102.166 Up 0 *
PE1-PE2-container-100-2
Total 2 displayed, Up 2, Down 0
```

After the ge-0/0/2 link between Routers P1 and P2 goes down, normalization is immediately triggered.

From operational mode, run the **show mpls container-lsp detail** command.

```
user@PE1> show mpls container-lsp detail
Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 4
Normalization
Min LSPs: 2, Max LSPs: 20
Aggregate bandwidth: 40Mbps, Sampled Aggregate bandwidth: 34.5538Mbps
NormalizeTimer: 3000 secs, NormalizeThreshold: 10%
Max Signaling BW: 10Mbps, Min Signaling BW: 10Mbps, Splitting BW: 40Mbps, Merging
BW: 6Mbps
Mode: incremental-normalization, failover-normalization
Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
Normalization in 2970 second(s)
11 Jun 5 19:28:27.564 Normalization complete: container (PE1-PE2-container-100)
with 4 members
10 Jun 5 19:28:20.315 Normalize: container (PE1-PE2-container-100) received PathErr
on member PE1-PE2-container-100-2[2 times]
9 Jun 5 19:28:20.315 Normalize: container (PE1-PE2-container-100) received
PathErr on member PE1-PE2-container-100-1[2 times]
8 Jun 5 19:28:20.311 Clear history and statistics: on container
(PE1-PE2-container-100)
7 Jun 5 19:28:20.311 Normalize: container (PE1-PE2-container-100) into 4
members - each with bandwidth 10000000 bps
6 Jun 5 19:28:20.311 Normalize: normalization with aggregate bandwidth 33665020
bps
5 Jun 5 19:28:20.308 Normalize: container (PE1-PE2-container-100) received
PathErr on member PE1-PE2-container-100-2
4 Jun 5 19:28:20.308 Normalize: container (PE1-PE2-container-100) received
PathErr on member PE1-PE2-container-100-1
3 Jun 5 19:27:48.574 Normalization complete: container (PE1-PE2-container-100)
with 2 members
2 Jun 5 19:27:28.644 Normalize: container (PE1-PE2-container-100) create 2
LSPs, min bw 10000000bps, member count 0
1 Jun 5 19:27:28.644 Normalize: normalization with aggregate bandwidth 0 bps
----Output truncated----
```

**Meaning** Arrival of PathErr message on link failure triggers immediate normalization.

### Verifying Incremental Normalization

**Purpose** Verify incremental normalization when enough bandwidth is not available.

On Router PE1, the RSVP interfaces static bandwidth is restricted to 22 Mbps each.

**Action** From operational mode, run the **show rsvp interface** command.

```
user@PE1> show rsvp interface
RSVP interface: 4 active
```

| Interface  | State | Active resv | Subscription | Static BW | Available BW | Reserved BW | Highwater mark |
|------------|-------|-------------|--------------|-----------|--------------|-------------|----------------|
| ge-0/0/2.0 | Up    | 0           | 100%         | 22Mbps    | 22Mbps       | 0bps        | 21.4031Mbps    |
| ge-0/0/1.0 | Up    | 2           | 100%         | 22Mbps    | 12Mbps       | 10Mbps      | 21.7011Mbps    |

Before normalization happens:

From operational mode, run the **show mpls container-lsp** command.

```
user@PE1> show mpls container-lsp
Ingress LSP: 1 sessions
Container LSP name
PE1-PE2-container-100
To From State Rt P State Up Member LSP count
ActivePath LSPname
10.255.102.128 10.255.102.166 Up 0 * PE1-PE2-container-100-1
PE1-PE2-container-100-1
10.255.102.128 10.255.102.166 Up 0 * PE1-PE2-container-100-2
```

After normalization happens:

From operational mode, run the **show mpls container-lsp** command.

```
user@PE1> show mpls container-lsp
Ingress LSP: 1 sessions
Container LSP name
PE1-PE2-container-100
To From State Rt P State Up Member LSP count
ActivePath LSPname
10.255.102.128 10.255.102.166 Up 0 * PE1-PE2-container-100-1
10.255.102.128 10.255.102.166 Up 0 * PE1-PE2-container-100-2
10.255.102.128 10.255.102.166 Up 0 * PE1-PE2-container-100-3
10.255.102.128 10.255.102.166 Up 0 * PE1-PE2-container-100-4
10.255.102.128 10.255.102.166 Up 0 * PE1-PE2-container-100-5
10.255.102.128 10.255.102.166 Up 0 * PE1-PE2-container-100-6
10.255.102.128 0.0.0.0 Dn 0 - PE1-PE2-container-100-7
Total 7 displayed, Up 6, Down 1
```

From operational mode, run the **show mpls container-lsp detail** command.

```
user@PE1> show mpls container-lsp detail
Ingress LSP: 1 sessions
Container LSP name: PE1-PE2-container-100, State: Up, Member count: 7
Normalization
Min LSPs: 2, Max LSPs: 10
Aggregate bandwidth: 40.8326Mbps, Sampled Aggregate bandwidth: 50.129Mbps
NormalizeTimer: 9000 secs, NormalizeThreshold: 10%
Max Signaling BW: 10Mbps, Min Signaling BW: 5Mbps, Splitting BW: 40Mbps, Merging BW: 5Mbps
Mode: incremental-normalization, failover-normalization
Sampling: Outlier cut-off 1, Percentile 90 of Aggregate
Normalization in 8072 second(s)
10 Jun 5 18:40:17.812 Normalization complete: container (PE1-PE2-container-100) with 7 members, retry-limit reached
9 Jun 5 18:40:08.028 Normalize: container (PE1-PE2-container-100) for target member count 7, member bandwidth 6805439 bps
8 Jun 5 18:39:58.301 Normalize: container (PE1-PE2-container-100) for target
```

```

member count 6, member bandwidth 7939679 bps
7 Jun 5 18:39:48.470 Clear history and statistics: on container
(PE1-PE2-container-100)
6 Jun 5 18:39:48.470 Normalize: container (PE1-PE2-container-100) into 5
members - each with bandwidth 9527615 bps
5 Jun 5 18:39:48.469 Normalize: normalization with aggregate bandwidth 47638076
bps
4 Jun 5 18:39:48.469 Normalize: normalization with 47638076 bps
3 Jun 5 18:39:09.471 Normalization complete: container (PE1-PE2-container-100)
with 2 members
2 Jun 5 18:38:59.822 Normalize: container (PE1-PE2-container-100) create 2
LSPs, min bw 5000000bps, member count 0
1 Jun 5 18:38:59.822 Normalize: normalization with aggregate bandwidth 0 bps

```

**Meaning** After normalization, the aggregate bandwidth after three retries is 40.8326 Mbps.

**Related Documentation** • [Dynamic Bandwidth Management Using Container LSP Overview on page 382](#)

## Configuring On-Demand Loss and Delay Measurement

You can configure an on-demand loss and delay measurement for point-to-point ultimate hop popping (UHP) label-switched paths (LSPs) in MPLS networks to monitor network performance. The **monitor mpls loss rsvp**, **monitor mpls delay rsvp**, and **monitor mpls loss-delay rsvp** CLI commands provide an on-demand summary of performance metrics for direct mode packet loss, two-way packet delay, and related metrics, such as inter-packet delay variation and channel throughput measurement.

This functionality provides real-time visibility into network performance, thereby facilitating network performance planning, troubleshooting, and evaluation.

Before you begin:

1. Configure the device interfaces.
2. Configure the device router ID.
3. Configure the following protocols:
  - RSVP
  - OSPF

Enable traffic engineering capabilities.

  - MPLS

To configure the PE device:

1. Enable the chassis with tunnel services and enhanced IP network services configuration.
 

```

[edit chassis]
user@R1# set fpc fpc-slot pic pic-slot tunnel-services bandwidth bandwidth
user@R1# set network-services enhanced-ip

```
2. Configure an associated bidirectional LSP to the remote router.

```
[edit protocols]
user@R1# set mpls label-switched-path lsp-name to remote-router-ip-address
user@R1# set mpls label-switched-path lsp-name oam mpls-tp-mode
user@R1# set mpls label-switched-path lsp-name ultimate-hop-popping
user@R1# set mpls label-switched-path lsp-name associate-lsp lsp-name
```

3. Create traffic classes for maintaining data traffic statistics per traffic class.

This enables traffic class scoped loss measurement.

```
[edit protocols]
user@R1# set mpls statistics traffic-class-statistics
```

#### Related Documentation

- [On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview on page 50](#)
- [Example: Configuring On-Demand Loss and Delay Measurement on page 439](#)
- [monitor mpls loss rsvp on page 1154](#)
- [monitor mpls delay rsvp on page 1150](#)
- [monitor mpls loss-delay rsvp on page 1159](#)

## Example: Configuring On-Demand Loss and Delay Measurement

This example shows how to enable on-demand loss and delay measurement for point-to-point ultimate hop popping (UHP) label-switched paths (LSPs) in MPLS networks to monitor network performance.

- [Requirements on page 439](#)
- [Overview on page 440](#)
- [Configuration on page 440](#)
- [Verification on page 444](#)

### Requirements

This example uses the following hardware and software components:

- Two MX Series 3D Universal Edge routers that contain MPC/MICs only
- Junos OS Release 14.2 or later running on all the routers

Before you begin:

1. Configure the device interfaces.
2. Configure the autonomous system numbers and router IDs for the devices.
3. Configure the following protocols:
  - RSVP
  - MPLS
  - OSPF

## Overview

Starting with Junos OS Release 14.2, an on-demand tool to monitor and measure packet loss, packet delay, or both for associated bidirectional MPLS ultimate hop popping (UHP) point-to-point label-switched paths (LSPs) is introduced. The tool can be enabled using the following CLI commands – **monitor mpls loss rsvp**, **monitor mpls delay rsvp**, and **monitor mpls loss-delay rsvp**.


These commands provide an on-demand summary of performance metrics for direct mode packet loss, two-way packet delay, and related metrics, such as inter-packet delay variation and channel throughput measurement.

This functionality provides real-time visibility into network performance, thereby facilitating network performance planning, troubleshooting, and evaluation.

## Topology

Figure 38 on page 440 illustrates the on-demand loss and delay measurement using a simple two-router topology.

**Figure 38: Configuring On-Demand Loss and Delay Measurement**



In this example, an associated bidirectional LSP is configured between Routers R1 and R2, for which the performance metrics is measured.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```

R1
set chassis fpc 0 pic 3 tunnel-services bandwidth 1g
set chassis network-services enhanced-ip
set interfaces ge-0/0/0 unit 0 family inet address 1.1.1.1/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.0.0.1/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 10.0.0.1
set protocols rsvp interface ge-0/0/0.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fxp0.0 disable
set protocols mpls statistics traffic-class-statistics
set protocols mpls label-switched-path R1-R2 to 20.0.0.1
set protocols mpls label-switched-path R1-R2 oam mpls-tp-mode
set protocols mpls label-switched-path R1-R2 ultimate-hop-popping
set protocols mpls label-switched-path R1-R2 associate-lsp R2-R1
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable

```

```

set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

```

R2 set chassis fpc 0 pic 3 tunnel-services bandwidth 1g
set chassis network-services enhanced-ip
set interfaces ge-0/0/0 unit 0 family inet address 1.1.1.2/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 20.0.0.1/32
set interfaces lo0 unit 0 family mpls
set routing-options router-id 20.0.0.1
set protocols rsvp interface ge-0/0/0.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fxp0.0 disable
set protocols mpls statistics traffic-class-statistics
set protocols mpls label-switched-path R2-R1 to 10.0.0.1
set protocols mpls label-switched-path R2-R1 oam mpls-tp-mode
set protocols mpls label-switched-path R2-R1 ultimate-hop-popping
set protocols mpls label-switched-path R2-R1 associate-lsp R1-R2
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router R1:

1. Enable the chassis with tunnel services and enhanced IP network services configuration.  

```

[edit chassis]
user@R1# set fpc 0 pic 3 tunnel-services bandwidth 1g
user@R1# set network-services enhanced-ip

```
2. Configure the interfaces for Router R1.  

```

[edit interfaces]
user@R1# set ge-0/0/0 unit 0 family inet address 1.1.1.1/30
user@R1# set ge-0/0/0 unit 0 family mpls

user@R1# set lo0 unit 0 family inet address 10.0.0.1/32
user@R1# set lo0 unit 0 family mpls

```
3. Configure the router ID for Router R1.  

```

[edit routing-options]
user@R1# set router-id 10.0.0.1

```
4. Enable RSVP on all the interfaces of Router R1, excluding the management interface.

```
[edit protocols]
user@R1# set rsvp interface ge-0/0/0.0
user@R1# set rsvp interface lo0.0
user@R1# set rsvp interface fxp0.0 disable
```

5. Enable MPLS on all the interfaces of Router R1, excluding the management interface.

```
[edit protocols]
user@R1# set mpls interface ge-0/0/0.0
user@R1# set mpls interface lo0.0
user@R1# set mpls interface fxp0.0 disable
```

6. Configure an associated bidirectional LSP to Router R2.

```
[edit protocols]
user@R1# set mpls label-switched-path R1-R2 to 20.0.0.1
user@R1# set mpls label-switched-path R1-R2 oam mpls-tp-mode
user@R1# set mpls label-switched-path R1-R2 ultimate-hop-popping
user@R1# set mpls label-switched-path R1-R2 associate-lsp R2-R1
```

7. Create traffic classes for maintaining data traffic statistics per traffic class.

This enables traffic class scoped loss measurement.

```
[edit protocols]
user@R1# set mpls statistics traffic-class-statistics
```

8. Configure OSPF with traffic engineering capabilities, and enable OSPF on all the interfaces of Router R1, excluding the management interface.

```
[edit protocols]
user@R1# set ospf traffic-engineering
user@R1# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@R1# set ospf area 0.0.0.0 interface lo0.0
user@R1# set ospf interface fxp0.0 disable
```

---

## Results

From configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces**, **show routing-options**, and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show chassis
fpc 0 {
 pic 3 {
 tunnel-services {
 bandwidth 1g;
 }
 }
}
network-services enhanced-ip;

user@R1# show interfaces
ge-0/0/0 {
 unit 0 {
 family inet {
 address 1.1.1.1/30;
 }
 }
}
```



```

 }
 family mpls;
 }
}
lo0 {
 unit 0 {
 family inet {
 address 10.0.0.1/32;
 }
 family mpls;
 }
}

user@R1# show routing-options
router-id 10.0.0.1;

user@R1# show protocols
rsvp {
 interface ge-0/0/0.0;
 interface lo0.0;
 interface fxp0.0 {
 disable;
 }
}
mpls {
 statistics {
 traffic-class-statistics;
 }
 label-switched-path R1-R2 {
 to 20.0.0.1;
 oam mpls-tp-mode;
 ultimate-hop-popping;
 associate-lsp R2-R1;
 }
 interface ge-0/0/0.0;
 interface lo0.0;
 interface fxp0.0 {
 disable;
 }
}
ospf {
 traffic-engineering;
 area 0.0.0.0 {
 interface ge-0/0/0.0;
 interface lo0.0;
 interface fxp0.0 {
 disable;
 }
 }
}
}

```

## Verification

Confirm that the configuration is working properly.

- [Verifying the LSP Status on page 444](#)
- [Verifying Packet Loss Measurement on page 444](#)
- [Verifying Packet Delay Measurement on page 446](#)
- [Verifying Packet Loss-Delay Measurement on page 446](#)

---

### Verifying the LSP Status

**Purpose** Verify that the associated bidirectional LSP between Routers R1 and R2 is up.

**Action** From operational mode, run the **show mpls lsp** command.

```
user@R1> show mpls lsp
Ingress LSP: 1 sessions
To From State Rt P ActivePath LSPname
20.0.0.1 10.0.0.1 Up 0 * R1-R2 Assoc-Bidir
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
To From State Rt Style Labelin Labelout LSPname
10.0.0.1 20.0.0.1 Up 0 1 FF 299776 - R2-R1 Assoc-Bidir
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

**Meaning** The associated bidirectional LSP R1-R2 is up and active.

---

### Verifying Packet Loss Measurement

**Purpose** Verify the on-demand loss measurement result.

**Action** From operational mode, run the **monitor mpls loss rsvp R1-R2 count 2 detail** command.

```

user@R1> monitor mpls loss rsvp R1-R2 count 2 detail
(0)
Response code : Success
Origin timestamp : 1404129082 secs, 905571890 nsecs
Forward transmit count : 83040
Forward receive count : 83040
Reverse transmit count : 83100
Reverse receive count : 83100
(1)
Response code : Success
Origin timestamp : 1404129083 secs, 905048410 nsecs
Forward transmit count : 83841
Forward receive count : 83841
Reverse transmit count : 83904
Reverse receive count : 83904
Current forward transmit count : 801
Current forward receive count : 801
Current forward loss : 0 packets
Current forward loss ratio : 0.000000
Current forward throughput : 0.801 kpps
Current reverse transmit count : 804
Current reverse receive count : 804
Current reverse loss : 0 packets
Current reverse loss ratio : 0.000000
Current reverse throughput : 0.804 kpps
(2)
Response code : Success
Origin timestamp : 1404129084 secs, 904828715 nsecs
Forward transmit count : 84423
Forward receive count : 84423
Reverse transmit count : 84487
Reverse receive count : 84487
Current forward transmit count : 582
Current forward receive count : 582
Current forward loss : 0 packets
Current forward loss ratio : 0.000000
Current forward throughput : 0.582 kpps
Current reverse transmit count : 583
Current reverse receive count : 583
Current reverse loss : 0 packets
Current reverse loss ratio : 0.000000
Current reverse throughput : 0.583 kpps

Cumulative forward transmit count : 1383
Cumulative forward loss : 0 packets
Average forward loss ratio : 0.000000
Average forward throughput : 0.692 kpps
Cumulative reverse transmit count : 1387
Cumulative reverse loss : 0 packets
Average reverse loss ratio : 0.000000
Average reverse throughput : 0.694 kpps

LM queries sent : 3
LM responses received : 3
LM queries timedout : 0
LM responses dropped due to errors : 0

```

**Meaning** The packet loss measurement for two counts is displayed.

### Verifying Packet Delay Measurement

---

**Purpose** Verify the on-demand delay measurement result.

**Action** From operational mode, run the **monitor mpls delay rsvp R1-R2 count 2 detail** command.

```

user@R1> monitor mpls delay rsvp R1-R2 count 2 detail
(1)
Response code : Success
Querier transmit timestamp : 1404129122 secs, 479955401 nsecs
Responder receive timestamp : 1404129122 secs, 468519022 nsecs
Responder transmit timestamp : 1404129122 secs, 470255123 nsecs
Querier receive timestamp : 1404129122 secs, 481736403 nsecs
Current two-way channel delay : 44 usecs
Current round-trip-time : 1781 usecs
(2)
Response code : Success
Querier transmit timestamp : 1404129123 secs, 480926210 nsecs
Responder receive timestamp : 1404129123 secs, 469488696 nsecs
Responder transmit timestamp : 1404129123 secs, 471130706 nsecs
Querier receive timestamp : 1404129123 secs, 482613911 nsecs
Current two-way channel delay : 45 usecs
Current round-trip-time : 1687 usecs

Best two-way channel delay : 44 usecs
Worst two-way channel delay : 45 usecs
Average two-way channel delay : 45 usecs
Best round-trip-time : 1687 usecs
Worst round-trip-time : 1781 usecs
Average round-trip-time : 1734 usecs
Average forward delay variation : 1 usecs
Average reverse delay variation : 1 usecs

DM queries sent : 2
DM responses received : 2
DM queries timedout : 0
DM responses dropped due to errors : 0

```

**Meaning** The packet delay measurement for two counts is displayed.

### Verifying Packet Loss-Delay Measurement

---

**Purpose** Verify the on-demand loss and delay measurement result.

**Action** From operational mode, run the **monitor mpls loss-delay rsvp R1-R2 count 2 detail** command.

```

user@R1> monitor mpls loss-delay rsvp R1-R2 count 2 detail
(0)
Response code : Success
Forward transmit count : 142049
Forward receive count : 142049
Reverse transmit count : 142167
Reverse receive count : 142167
Querier transmit timestamp : 1404129161 secs, 554422723 nsecs
Responder receive timestamp : 1404129161 secs, 542877570 nsecs
Responder transmit timestamp : 1404129161 secs, 546004545 nsecs
Querier receive timestamp : 1404129161 secs, 557599327 nsecs
(1)
Response code : Success
Forward transmit count : 143049
Forward receive count : 143049
Reverse transmit count : 143168
Reverse receive count : 143168
Current forward transmit count : 1000
Current forward receive count : 1000
Current forward loss : 0 packets
Current forward loss ratio : 0.000000
Current forward throughput : 1.000 kpps
Current reverse transmit count : 1001
Current reverse receive count : 1001
Current reverse loss : 0 packets
Current reverse loss ratio : 0.000000
Current reverse throughput : 1.001 kpps
Querier transmit timestamp : 1404129162 secs, 554465742 nsecs
Responder receive timestamp : 1404129162 secs, 542919166 nsecs
Responder transmit timestamp : 1404129162 secs, 545812736 nsecs
Querier receive timestamp : 1404129162 secs, 557409175 nsecs
Current two-way channel delay : 49 usecs
Current round-trip-time : 2943 usecs
(2)
Response code : Success
Forward transmit count : 143677
Forward receive count : 143677
Reverse transmit count : 143799
Reverse receive count : 143799
Current forward transmit count : 628
Current forward receive count : 628
Current forward loss : 0 packets
Current forward loss ratio : 0.000000
Current forward throughput : 0.627 kpps
Current reverse transmit count : 631
Current reverse receive count : 631
Current reverse loss : 0 packets
Current reverse loss ratio : 0.000000
Current reverse throughput : 0.630 kpps
Querier transmit timestamp : 1404129163 secs, 556698575 nsecs
Responder receive timestamp : 1404129163 secs, 545150128 nsecs
Responder transmit timestamp : 1404129163 secs, 546918408 nsecs
Querier receive timestamp : 1404129163 secs, 558515047 nsecs
Current two-way channel delay : 48 usecs
Current round-trip-time : 1816 usecs

Cumulative forward transmit count : 1628

```

|                                     |              |
|-------------------------------------|--------------|
| Cumulative forward loss             | : 0 packets  |
| Average forward loss ratio          | : 0.000000   |
| Average forward throughput          | : 0.813 kpps |
| Cumulative reverse transmit count   | : 1632       |
| Cumulative reverse loss             | : 0 packets  |
| Average reverse loss ratio          | : 0.000000   |
| Average reverse throughput          | : 0.815 kpps |
| Best two-way channel delay          | : 48 usecs   |
| Worst two-way channel delay         | : 49 usecs   |
| Average two-way channel delay       | : 49 usecs   |
| Best round-trip-time                | : 1816 usecs |
| Worst round-trip-time               | : 3176 usecs |
| Average round-trip-time             | : 2645 usecs |
| Average forward delay variation     | : 1 usecs    |
| Average reverse delay variation     | : 0 usecs    |
| LDM queries sent                    | : 3          |
| LDM responses received              | : 3          |
| LDM queries timedout                | : 0          |
| LDM responses dropped due to errors | : 0          |

**Meaning** The packet loss and delay measurement for two counts is displayed.

- Related Documentation**
- [On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview on page 50](#)
  - [monitor mpls loss rsvp on page 1154](#)
  - [monitor mpls delay rsvp on page 1150](#)
  - [monitor mpls loss-delay rsvp on page 1159](#)

---

## Configuring Pro-Active Loss and Delay Measurements

You can configure pro-active loss and delay measurements for point-to-point ultimate-hop popping label-switched paths (LSPs) in MPLS networks to monitor network performance. The **show performance-monitoring mpls lsp** CLI command provides a summary of performance metrics for direct mode packet loss, two-way packet delay, and related metrics, such as inter-packet delay variation and channel throughput measurement.

This functionality provides real-time visibility into network performance, thereby facilitating network performance planning, troubleshooting, and evaluation.

This feature provides the following performance metrics:

- Inter-packet delay variation (IPDV)
- Loss measurement
- Round-trip delay (RTT)
- Throughput measurement
- Two-way channel delay

Before you begin:

1. Configure the device interfaces.
2. Configure the autonomous system numbers and router IDs for the devices.
3. Configure the following protocols:
  - MPLS
  - OSPF
  - RSVP

To configure pro-active loss and delay measurements on the PE device:

1. Configure an associated bidirectional LSP to Router R2.

```
[edit protocols]
user@host# set mpls label-switched-path lsp-name to remote-router-ip-address
user@host# set mpls label-switched-path lsp-name install
destination-prefix/prefix-length active
user@host# set mpls label-switched-path lsp-name oam mpls-tp-mode
user@host# set mpls label-switched-path lsp-name ultimate-hop-popping
user@host# set mpls label-switched-path lsp-name associate-lsp remote-lsp-name
```

2. Create traffic classes for maintaining data traffic statistics per traffic class.

This enables traffic class scoped loss and delay measurements.

```
[edit protocols]
user@host# set mpls statistics traffic-class-statistics
```

3. Configure performance monitoring at the querier side.

```
[edit protocols]
user@host# set mpls label-switched-path lsp-name oam performance-monitoring
querier delay traffic-class tc-value query-interval milliseconds
user@host# set mpls label-switched-path lsp-name oam performance-monitoring
querier loss traffic-class tc-value query-interval milliseconds
user@host# set mpls label-switched-path lsp-name oam performance-monitoring
querier loss-delay traffic-class tc-value query-interval milliseconds
```

4. Configure performance monitoring at the responder side.

```
[edit protocols]
user@host# set mpls label-switched-path lsp-name oam performance-monitoring
responder delay min-query-interval milliseconds
user@host# set mpls label-switched-path lsp-name oam performance-monitoring
responder loss min-query-interval milliseconds
```

#### Related Documentation

- [Example: Configuring Pro-active Loss and Delay Measurements for Bidirectional MPLS LSPs on page 450](#)
- [performance-monitoring \(Protocols MPLS\) on page 922](#)
- [show performance-monitoring mpls lsp on page 1236](#)

## Example: Configuring Pro-active Loss and Delay Measurements for Bidirectional MPLS LSPs

---

This example shows how to configure pro-active loss and delay measurements for point-to-point ultimate-hop popping label-switched paths (LSPs) in MPLS networks to monitor network performance.

- [Requirements on page 450](#)
- [Overview on page 450](#)
- [Configuration on page 451](#)
- [Verification on page 455](#)

### Requirements

This example uses the following hardware and software components:

- Two MX Series 3D Universal Edge routers that contain MPC/MICs only
- Junos OS Release 15.1 or later running on all the routers

Before you begin:

1. Configure the device interfaces.
2. Configure the autonomous system numbers and router IDs for the devices.
3. Configure the following protocols:
  - a. MPLS
  - b. OSPF
  - c. RSVP

### Overview

Starting with Junos OS Release 15.1, a pro-active tool to monitor and measure packet loss, packet delay, or both for associated bidirectional MPLS ultimate-hop popping point-to-point label-switched paths (LSPs) is introduced.

This feature provides the following performance metrics:

- Inter-packet delay variation (IPDV)
- Loss measurement
- Round-trip delay (RTT)
- Throughput measurement
- Two-way channel delay


This functionality provides real-time visibility into network performance, thereby facilitating network performance planning, troubleshooting, and evaluation.



## Topology

Figure 38 on page 440 illustrates the pro-active loss and delay measurements using a simple two-router topology.

Figure 39: Configuring Pro-Active Loss and Delay Measurements



In this example, an associated bidirectional LSP is configured between Routers R1 and R2, for which the performance metrics are measured.

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```

R1 set chassis network-services enhanced-ip
 set interfaces ge-0/0/0 unit 0 family inet address 1.1.1.1/30
 set interfaces ge-0/0/0 unit 0 family mpls
 set interfaces lo0 unit 0 family inet address 10.0.0.1/32
 set interfaces lo0 unit 0 family mpls
 set protocols mpls interface ge-0/0/0.0
 set protocols mpls interface lo0.0
 set protocols mpls interface fxp0.0 disable
 set protocols mpls label-switched-path R1-R2 associate-lsp R2-R1
 set protocols mpls label-switched-path R1-R2 install 20.10.30.0/24 active
 set protocols mpls label-switched-path R1-R2 oam mpls-tp-mode
 set protocols mpls label-switched-path R1-R2 oam performance-monitoring querier delay
 traffic-class tc-0 query-interval 1000
 set protocols mpls label-switched-path R1-R2 oam performance-monitoring querier loss
 traffic-class none query-interval 1000
 set protocols mpls label-switched-path R1-R2 oam performance-monitoring querier
 loss-delay traffic-class tc-0 query-interval 1000
 set protocols mpls label-switched-path R1-R2 oam performance-monitoring responder
 delay min-query-interval 1000
 set protocols mpls label-switched-path R1-R2 oam performance-monitoring responder
 loss min-query-interval 1000
 set protocols mpls label-switched-path R1-R2 to 20.0.0.1
 set protocols mpls label-switched-path R1-R2 ultimate-hop-popping
 set protocols mpls statistics traffic-class-statistics
 set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
 set protocols ospf area 0.0.0.0 interface lo0.0
 set protocols ospf area 0.0.0.0 interface fxp0.0 disable
 set protocols ospf traffic-engineering
 set protocols rsvp interface ge-0/0/0.0
 set protocols rsvp interface lo0.0
 set protocols rsvp interface fxp0.0 disable
 set routing-options router-id 10.0.0.1

R2 set chassis network-services enhanced-ip

```

```

set interfaces ge-0/0/0 unit 0 family inet address 1.1.1.2/30
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 20.0.0.1/32
set interfaces lo0 unit 0 family mpls
set protocols mpls interface ge-0/0/0.0
set protocols mpls interface lo0.0
set protocols mpls interface fxp0.0 disable
set protocols mpls label-switched-path R2-R1 associate-lsp R1-R2
set protocols mpls label-switched-path R2-R1 install 10.10.20.0/24 active
set protocols mpls label-switched-path R2-R1 oam mpls-tp-mode
set protocols mpls label-switched-path R2-R1 oam performance-monitoring responder
 delay min-query-interval 1000
set protocols mpls label-switched-path R2-R1 oam performance-monitoring responder
 loss min-query-interval 1000
set protocols mpls label-switched-path R2-R1 oam performance-monitoring querier delay
 traffic-class tc-0 query-interval 1000
set protocols mpls label-switched-path R2-R1 oam performance-monitoring querier loss
 traffic-class none query-interval 1000
set protocols mpls label-switched-path R2-R1 oam performance-monitoring querier
 loss-delay traffic-class tc-0 query-interval 1000
set protocols mpls label-switched-path R2-R1 to 10.0.0.1
set protocols mpls label-switched-path R2-R1 ultimate-hop-popping
set protocols mpls statistics traffic-class-statistics
set protocols ospf area 0.0.0.0 interface ge-0/0/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols rsvp interface ge-0/0/0.0
set protocols rsvp interface lo0.0
set protocols rsvp interface fxp0.0 disable
set routing-options router-id 20.0.0.1

```

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router R1:

1. Enable the enhanced IP network services configuration.
 

```

[edit chassis]
user@R1# set network-services enhanced-ip

```
2. Configure the interfaces for Router R1.
 

```

[edit interfaces]
user@R1# set ge-0/0/0 unit 0 family inet address 1.1.1.1/30
user@R1# set ge-0/0/0 unit 0 family mpls

user@R1# set lo0 unit 0 family inet address 10.0.0.1/32
user@R1# set lo0 unit 0 family mpls

```
3. Configure the router ID for Router R1.
 

```

[edit routing-options]
user@R1# set router-id 10.0.0.1

```

4. Enable RSVP on all the interfaces of Router R1, excluding the management interface.
 

```
[edit protocols]
user@R1# set rsvp interface ge-0/0/0.0
user@R1# set rsvp interface lo0.0
user@R1# set rsvp interface fxp0.0 disable
```
5. Enable MPLS on all the interfaces of Router R1, excluding the management interface.
 

```
[edit protocols]
user@R1# set mpls interface ge-0/0/0.0
user@R1# set mpls interface lo0.0
user@R1# set mpls interface fxp0.0 disable
```
6. Configure an associated bidirectional LSP to Router R2.
 

```
[edit protocols]
user@R1# set mpls label-switched-path R1-R2 to 20.0.0.1
user@R1# set mpls label-switched-path R1-R2 install 20.10.30.0/24 active
user@R1# set mpls label-switched-path R1-R2 oam mpls-tp-mode
user@R1# set mpls label-switched-path R1-R2 ultimate-hop-popping
user@R1# set mpls label-switched-path R1-R2 associate-lsp R2-R1
```
7. Create traffic classes for maintaining data traffic statistics per traffic class.
 

This enables traffic class scoped loss and delay measurement.

```
[edit protocols]
user@R1# set mpls statistics traffic-class-statistics
```
8. Configure performance monitoring at the querier side.
 

```
[edit protocols]
user@R1# set mpls label-switched-path R1-R2 oam performance-monitoring querier
 delay traffic-class tc-0 query-interval 1000
user@R1# set mpls label-switched-path R1-R2 oam performance-monitoring querier
 loss traffic-class none query-interval 1000
user@R1# set mpls label-switched-path R1-R2 oam performance-monitoring querier
 loss-delay traffic-class tc-0 query-interval 1000
```
9. Configure performance monitoring at the responder side.
 

```
[edit protocols]
user@R1# set mpls label-switched-path R1-R2 oam performance-monitoring
 responder delay min-query-interval 1000
user@R1# set mpls label-switched-path R1-R2 oam performance-monitoring
 responder loss min-query-interval 1000
```
10. Configure OSPF with traffic engineering capabilities, and enable OSPF on all the interfaces of Router R1, excluding the management interface.
 

```
[edit protocols]
user@R1# set ospf traffic-engineering
user@R1# set ospf area 0.0.0.0 interface ge-0/0/0.0
user@R1# set ospf area 0.0.0.0 interface lo0.0
user@R1# set ospf interface fxp0.0 disable
```

## Results

---

From configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces**, **show routing-options**, and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show chassis
network-services enhanced-ip;

user@R1# show interfaces
ge-0/0/0 {
 unit 0 {
 family inet {
 address 1.1.1.1/30;
 }
 family mpls;
 }
}
lo0 {
 unit 0 {
 family inet {
 address 10.0.0.1/32;
 }
 family mpls;
 }
}

user@R1# show routing-options
router-id 10.0.0.1;

user@R1# show protocols
rsvp {
 interface ge-0/0/0.0;
 interface lo0.0;
 interface fxp0.0 {
 disable;
 }
}
mpls {
 label-switched-path R1-R2 {
 to 20.0.0.1;
 install 20.10.30.0/24 active;
 oam {
 mpls-tp-mode;
 performance-monitoring {
 querier {
 loss {
 traffic-class none {
 query-interval 1000;
 }
 }
 }
 }
 delay {
 traffic-class tc-0 {
 query-interval 1000;
 }
 }
 }
 }
}
```

```

 }
 loss-delay {
 traffic-class none {
 query-interval 1000;
 }
 }
}
responder {
 loss {
 min-query-interval 1000;
 }
 delay {
 min-query-interval 1000;
 }
}
}
ultimate-hop-popping;
associate-lsp R2-R1;
}
}
ospf {
 traffic-engineering;
 area 0.0.0.0 {
 interface ge-0/0/0.0;
 interface lo0.0;
 interface fxp0.0 {
 disable;
 }
 }
}
}

```

## Verification

### Verifying Loss and Delay Measurement

**Purpose** Verify the loss and delay measurement.

**Action** From operational mode, run the **show performance-monitoring mpls lsp** command.

```
user@R1> show performance-monitoring mpls lsp
Session Total: 3 Up: 3 Down: 0
LSP name:R1-R2, PM State:Up
 Loss measurement Data:
 Duration: 00:04:43
 Traffic-class: None
 Queries sent: 282
 Responses received: 282
 Responses dropped due to errors: 0
 Queries timeout: 0
 Forward loss measurement:
 Average packet loss: 0
 Average packet throughput: 554338
 Reverse loss measurement:
 Average packet loss: 0
 Average packet throughput: 1352077
LSP name:R1-R2, PM State:Up
 Delay measurement Data:
 Duration: 00:04:43
 Traffic-class: 0
 Queries sent: 282
 Responses received: 282
 Responses dropped due to errors: 0
 Queries timeout: 0
 Best 2-way channel delay: 72 usecs
 Worst 2-way channel delay: 365 usecs
 Best round trip time: 843 usecs
 Worst round trip time: 105523 usecs
 Avg absolute fw delay variation: 1619 usecs
 Avg absolute rv delay variation: 1619 usecs
LSP name:R1-R2, PM State:Up
 Loss measurement Data:
 Duration: 00:04:43
 Traffic-class: None
 Queries sent: 282
 Responses received: 282
 Responses dropped due to errors: 0
 Queries timeout: 0
 Forward loss measurement:
 Average packet loss: 0
 Average packet throughput: 553927
 Reverse loss measurement:
 Average packet loss: 0
 Average packet throughput: 1351531
 Delay measurement Data:
 Best 2-way channel delay: 76 usecs
 Worst 2-way channel delay: 368 usecs
 Best round trip time: 1082 usecs
 Worst round trip time: 126146 usecs
 Avg absolute fw delay variation: 1618 usecs
 Avg absolute rv delay variation: 1619 usecs
```

**Meaning** The packet loss and delay measurement metrics for LSP are displayed.

**Related Documentation**

- [performance-monitoring \(Protocols MPLS\) on page 922](#)
- [show performance-monitoring mpls lsp on page 1236](#)

## PART 3

# Configuring RSVP

- [RSVP Overview on page 459](#)
- [Configuring RSVP on page 471](#)
- [Configuring RSVP Link Protection and Node Protection to Protect from Traffic Failures on page 495](#)
- [Configuring RSVP Graceful Restart for High Availability on page 511](#)





## CHAPTER 9

# RSVP Overview

- [RSVP Overview on page 459](#)
- [Supported RSVP Standards on page 460](#)
- [Junos OS RSVP Protocol Implementation on page 461](#)
- [RSVP Operation Overview on page 462](#)
- [RSVP Authentication on page 462](#)
- [RSVP and IGP Hello Packets and Timers on page 462](#)
- [RSVP Message Types on page 463](#)
- [Understanding RSVP Automatic Mesh on page 463](#)
- [Path Messages on page 465](#)
- [Resv Messages on page 465](#)
- [PathTear Messages on page 465](#)
- [ResvTear Messages on page 465](#)
- [PathErr Messages on page 466](#)
- [ResvErr Messages on page 466](#)
- [ResvConfirm Messages on page 466](#)
- [RSVP Reservation Styles on page 466](#)
- [RSVP Refresh Reduction on page 467](#)
- [MTU Signaling in RSVP on page 468](#)
- [How the Correct MTU Is Signaled in RSVP on page 469](#)
- [Determining an Outgoing MTU Value on page 469](#)
- [MTU Signaling in RSVP Limitations on page 470](#)

## RSVP Overview

---

RSVP is a resource reservation setup protocol that is used by both network hosts and routers. Hosts use RSVP to request a specific class of service (CoS) from the network for particular application flows. Routers use RSVP to deliver CoS requests to all routers along the data path. RSVP also can maintain and refresh states for a requested CoS application flow.

RSVP treats an application flow as a simplex connection. That is, the CoS request travels only in one direction—from the sender to the receiver. RSVP is a transport layer protocol that uses IP as its network layer. However, RSVP does not transport application flows. Rather, it is more of an Internet control protocol, similar to the Internet Control Message Protocol (ICMP) and Internet Group Management Protocol (IGMP). RSVP runs as a separate software process in the Junos OS and is not in the packet forwarding path.

RSVP is not a routing protocol, but rather is designed to operate with current and future unicast and multicast routing protocols. The routing protocols are responsible for choosing the routes to use to forward packets, and RSVP consults local routing tables to obtain routes. RSVP only ensures the CoS of packets traveling along a data path.

The receiver in an application flow requests the preferred CoS from the sender. To do this, the receiver issues an RSVP CoS request on behalf of the local application. The request propagates to all routers in reverse direction of the data paths toward the sender. In this process, RSVP requests might be merged, resulting in a protocol that scales well when there are a large number of receivers.

Because the number of receivers in an application flow is likely to change and the flow of delivery paths might change during the life of an application flow, RSVP takes a soft-state approach in its design, creating and removing the protocol states in routers and hosts incrementally over time. RSVP sends periodic refresh messages to maintain its state and to recover from occasional lost messages. In the absence of refresh messages, RSVP states automatically time out and are deleted.

---

## Supported RSVP Standards

Junos OS substantially supports the following RFCs and Internet drafts, which define standards for RSVP.

- RFC 2205, *Resource ReSerVation Protocol (RSVP)—Version 1 Functional Specification*
- RFC 2210, *The Use of RSVP with IETF Integrated Services*
- RFC 2211, *Specification of the Controlled-Load Network Element Service*
- RFC 2212, *Specification of Guaranteed Quality of Service*
- RFC 2215, *General Characterization Parameters for Integrated Service Network Elements*
- RFC 2745, *RSVP Diagnostic Messages*
- RFC 2747, *RSVP Cryptographic Authentication* (updated by RFC 3097)
- 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*

The Null Service Object for maximum transmission unit (MTU) signaling in RSVP is not supported.

- RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions*

Only Section 9, “Fault Handling,” is supported.

- RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*
- RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*  
(OSPF extensions can carry traffic engineering information over unnumbered links.)
- RFC 4558, *Node-ID Based Resource Reservation Protocol (RSVP) Hello: A Clarification Statement*
- RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*  
The RRO node ID subobject is for use in inter-AS link and node protection configurations.
- RFC 4875, *Extensions to RSVP-TE for Point-to-Multipoint TE LSPs*

The following RFCs do not define standards, but provide information about RSVP and related technologies. The IETF classifies them variously as “Experimental” or “Informational.”

- RFC 2209, *Resource ReSerVation Protocol (RSVP)—Version 1 Message Processing Rules*
- RFC 2216, *Network Element Service Specification Template*
- 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*

#### Related Documentation

- [Supported GMPLS Standards on page 673](#)
- [Supported LDP Standards on page 520](#)
- [Supported MPLS Standards on page 20](#)
- [Accessing Standards Documents on the Internet](#)

## Junos OS RSVP Protocol Implementation

The Junos implementation of RSVP supports RSVP version 1. The software includes support for all mandatory objects and RSVP message types, and supports message integrity and node authentications through the Integrity object.

The primary purpose of the Junos RSVP software is to support dynamic signaling within MPLS label-switched paths (LSPs). Supporting resource reservations over the Internet is only a secondary purpose of the Junos OS implementation. Since supporting resource reservations is secondary, the Junos RSVP software does not support the following features:

- IP multicasting sessions.

- Traffic control. The software cannot make resource reservations for real-time video or audio sessions.

With regard to the protocol mechanism, packet processing, and RSVP objects supported, the Junos OS implementation of the software is interoperable with other RSVP implementations.

## RSVP Operation Overview

---

RSVP creates independent sessions to handle each data flow. A session is identified by a combination of the destination address, an optional destination port, and a protocol. Within a session, there can be one or more senders. Each sender is identified by a combination of its source address and source port. An out-of-band mechanism, such as a session announcement protocol or human communication, is used to communicate the session identifier to all senders and receivers.

A typical RSVP session involves the following sequence of events:

1. A potential sender starts sending RSVP path messages to the session address.
2. A receiver, wanting to join the session, registers itself if necessary. For example, a receiver in a multicast application would register itself with IGMP.
3. The receiver receives the path messages.
4. The receiver sends appropriate Resv messages toward the sender. These messages carry a flow descriptor, which is used by routers along the path to make reservations in their link-layer media.
5. The sender receives the Resv message and then starts sending application data.

This sequence of events is not necessarily strictly synchronized. For example, receivers can register themselves before receiving path messages from the sender, and application data can flow before the sender receives Resv messages. Application data that is delivered before the actual reservation contained in the Resv message typically is treated as best-effort, non-real-time traffic with no CoS guarantee.

## RSVP Authentication

---

The Junos OS supports both the RSVP authentication style described in RFC 2747 (allowing for multivendor compatibility) and the RSVP authentication style described in Internet draft draft-ietf-rsvp-md5-03.txt. The Junos OS uses the authentication style described in Internet draft draft-ietf-rsvp-md5-08.txt by default. If the router receives an RFC 2747-style RSVP authentication from a neighbor, it switches to this style of authentication for that neighbor. The RSVP authentication style for each neighboring router is determined separately.

## RSVP and IGP Hello Packets and Timers

---

RSVP monitors the status of the interior gateway protocol (IGP) (IS-IS or OSPF) neighbors and relies on the IGP protocols to detect when a node fails. If an IGP protocol declares a neighbor down (because hello packets are no longer being received), RSVP also brings

down that neighbor. However, the IGP protocols and RSVP still act independently when bringing a neighbor up.

In the Junos OS, RSVP typically relies on IGP hello packet detection to check for node failures. RSVP sessions are kept up even if RSVP hello packets are no longer being received, so long as the router continues to receive IGP hello packets. RSVP sessions are maintained until either the router stops receiving IGP hello packets or the RSVP Path and Resv messages time out. Configuring a short time for the IS-IS or OSPF hello timers allows these protocols to detect node failures quickly.

RSVP hellos can be relied on when the IGP does not recognize a particular neighbor (for example, if IGP is not enabled on the interface) or if the IGP is RIP (not IS-IS or OSPF). Also, the equipment of other vendors might be configured to monitor RSVP sessions based on RSVP hello packets. This equipment might also take an RSVP session down due to a loss of RSVP hello packets.

We do not recommend configuring a short RSVP hello timer. If quick discovery of a failed neighbor is needed, configure short IGP (OSPF or IS-IS) hello timers.

OSPF and IS-IS have infrastructure to manage rapid hello message sending and receiving reliably, even if the routing protocols or some other process are straining the processing capability of the router. Under the same circumstances, RSVP hellos might time out prematurely even though the neighbor is functioning normally.

---

## RSVP Message Types

RSVP uses the following types of messages to establish and remove paths for data flows, establish and remove reservation information, confirm the establishment of reservations, and report errors:

- [Path Messages on page 465](#)
- [Resv Messages on page 465](#)
- [PathTear Messages on page 465](#)
- [ResvTear Messages on page 465](#)
- [PathErr Messages on page 466](#)
- [ResvErr Messages on page 466](#)
- [ResvConfirm Messages on page 466](#)

---

## Understanding RSVP Automatic Mesh

When adding sites to BGP and MPLS VPNs that use RSVP signaling, more configuration is needed to add provider edge (PE) routers than is needed to add customer edge (CE) devices. RSVP automatic mesh helps to reduce this configuration burden.

Service providers often use BGP and MPLS VPNs to efficiently scale the network while delivering revenue-generating services. In these environments, BGP is used to distribute the VPN routing information across the service provider's network, while MPLS is used

to forward that VPN traffic from one VPN site to another. BGP and MPLS VPNs are based on a peer model. To add a new CE device (site) to an existing VPN, you need to configure the CE router at the new site and the PE router connected to the CE router. You do not have to modify the configuration of all of the other PE routers participating in the VPN. The other PE routers automatically learn about the routes associated with the new site, a process called automatic discovery (AD).

The requirements are a bit different if you need to add a new PE router to the network. A BGP and MPLS VPN requires that the BGP session be fully meshed and that there also be a full mesh of PE router-to-PE router MPLS label-switched paths (LSPs) between all of the PE routers in the network. When you add a new PE router to the network, all of the existing PE routers must be reconfigured to peer with the new PE router. Much of the configuration effort can be reduced if you configure BGP route reflectors (mitigating the full mesh requirement for BGP) and if you configure (LDP) as the signaling protocol for MPLS.

However, if you need to add a new PE router to a network configured with a full mesh of RSVP-signaled LSPs, you must reconfigure each of the PE routers to have a peer relationship with the new PE router. You can configure RSVP automatic mesh to address this particular operational scenario. When you enable RSVP automatic mesh, RSVP LSPs are dynamically created between a new PE router and the existing PE routers, eliminating the need to reconfigure all of the PE routers manually. For dynamic LSP creation to function properly, BGP must be configured to exchange routes between all of the participating PE routers. If two BGP peers do not exchange routes, it is not possible to configure a dynamic LSP between them. The local router's inet.0 routing table must include a labeled route to each potential IBGP next-hop (future potential PE routers or LSP destinations).

RSVP includes numerous capabilities that are not available in LDP, including fast reroute, end-point control, and link management. RSVP automatic mesh helps to reduce the operation and maintenance requirements for RSVP, making it possible to deploy RSVP in larger and more complicated networks.

Every PE router can reach every other PE router in the network because this information is distributed by the IGP. A PE router can set up a point-to-point RSVP LSP to any other PE router in the network as long as it knows that such an LSP is required. To build a full mesh of LSPs between the PE routers requires that each PE router know which of the other PE routers make up the full mesh.



**NOTE:** In Junos OS, RSVP automatic mesh is configured using the `rsvp-te` configuration statement at the `[edit routing-options dynamic-tunnels dynamic-tunnel-name]` hierarchy level. The `rsvp-te` configuration statement is also available for use in routing instances as a provider-tunnel option. When implemented as a provider-tunnel option, `rsvp-te` is used to configure the RSVP point-to-multipoint LSPs for multiprotocol BGP multicast VPNs, not to configure RSVP automatic mesh.

---

- Related Documentation**
- [Configuring RSVP Automatic Mesh on page 482](#)
  - *Example: Configuring RSVP Automatic Mesh*
  - [label-switched-path-template \(Multicast\) on page 990](#)

## Path Messages

Each sender host transmits path messages downstream along the routes provided by the unicast and multicast routing protocols. Path messages follow the exact paths of application data, creating path states in the routers along the way, thus enabling routers to learn the previous-hop and next-hop node for the session. Path messages are sent periodically to refresh path states.

The refresh interval is controlled by a variable called the **refresh-time**, which is the periodical refresh timer expressed in seconds. A path state times out if a router does not receive a specified number of consecutive path messages. This number is specified by a variable called **keep-multiplier**. Path states are kept for  $(\text{keep-multiplier} + 0.5) \times 1.5 \times \text{refresh-time}$  seconds.

## Resv Messages

Each receiver host sends reservation request (Resv) messages upstream toward senders and sender applications. Resv messages must follow exactly the reverse path of path messages. Resv messages create and maintain a reservation state in each router along the way.

Resv messages are sent periodically to refresh reservation states. The refresh interval is controlled by the same refresh time variable, and reservation states are kept for  $(\text{keep-multiplier} + 0.5) \times 1.5 \times \text{refresh-time}$  seconds.

## PathTear Messages

PathTear messages remove (tear down) path states as well as dependent reservation states in any routers along a path. PathTear messages follow the same path as path messages. A PathTear typically is initiated by a sender application or by a router when its path state times out.

PathTear messages are not required, but they enhance network performance because they release network resources quickly. If PathTear messages are lost or not generated, path states eventually time out when they are not refreshed, and the resources associated with the path are released.

## ResvTear Messages

ResvTear messages remove reservation states along a path. These messages travel upstream toward senders of the session. In a sense, ResvTear messages are the reverse of Resv messages. ResvTear messages typically are initiated by a receiver application or by a router when its reservation state times out.

ResvTear messages are not required, but they enhance network performance because they release network resources quickly. If ResvTear messages are lost or not generated, reservation states eventually time out when they are not refreshed, and the resources associated with the reservation are released.

## PathErr Messages

---

When path errors occur (usually because of parameter problems in a path message), the router sends a unicast PathErr message to the sender that issued the path message. PathErr messages are advisory; these messages do not alter any path state along the way.

## ResvErr Messages

---

When a reservation request fails, a ResvErr error message is delivered to all the receivers involved. ResvErr messages are advisory; these messages do not alter any reservation state along the way.

## ResvConfirm Messages

---

Receivers can request confirmation of a reservation request, and this confirmation is sent with a ResvConfirm message. Because of the complex RSVP flow-merging rules, a confirmation message does not necessarily provide end-to-end confirmation of the entire path. Therefore, ResvConfirm messages are an indication, not a guarantee, of potential success.

Juniper Networks routers never request confirmation using the ResvConfirm message; however, a Juniper Networks router can send a ResvConfirm message if it receives a request from another vendor's equipment.

## RSVP Reservation Styles

---

A reservation request includes options for specifying the reservation style. The reservation styles define how reservations for different senders within the same session are treated and how senders are selected.

Two options specify how reservations for different senders within the same session are treated:

- Distinct reservation—Each receiver establishes its own reservation with each upstream sender.
- Shared reservation—All receivers make a single reservation that is shared among many senders.

Two options specify how senders are selected:

- Explicit sender—List all selected senders.
- Wildcard sender—Select all senders, which then participate in the session.



The following reservation styles, formed by a combination of these four options, currently are defined:

- **Fixed filter (FF)**—This reservation style consists of distinct reservations among explicit senders. Examples of applications that use fixed-filter-style reservations are video applications and unicast applications, which both require flows that have a separate reservation for each sender. The fixed filter reservation style is enabled on RSVP LSPs by default.
- **Wildcard filter (WF)**—This reservation style consists of shared reservations among wildcard senders. This type of reservation reserves bandwidth for any and all senders, and propagates upstream toward all senders, automatically extending to new senders as they appear. A sample application for wildcard filter reservations is an audio application in which each sender transmits a distinct data stream. Typically, only a few senders are transmitting at any one time. Such a flow does not require a separate reservation for each sender; a single reservation is sufficient.
- **Shared explicit (SE)**—This reservation style consists of shared reservations among explicit senders. This type of reservation reserves bandwidth for a limited group of senders. A sample application is an audio application similar to that described for wildcard filter reservations.

## RSVP Refresh Reduction

RSVP relies on soft-state to maintain the path and reservation states on each router. If the corresponding refresh messages are not sent periodically, the states eventually time out and reservations are deleted. RSVP also sends its control messages as IP datagrams with no reliability guarantee. It relies on periodic refresh messages to handle the occasional loss of Path or Resv messages.

The RSVP refresh reduction extensions, based on RFC 2961, addresses the following problems that result from relying on periodic refresh messages to handle message loss:

- **Scalability**—The scaling problem arises from the periodic transmission and processing overhead of refresh messages, which increases as the number of RSVP sessions increases.
- **Reliability and latency**—The reliability and latency problem stems from the loss of nonrefresh RSVP messages or one-time RSVP messages such as PathTear or PathErr. The time to recover from such a loss is usually tied to refresh interval and the keepalive timer.

The RSVP refresh reduction capability is advertised by enabling the refresh reduction (RR) capable bit in the RSVP common header. This bit is only significant between RSVP neighbors.

RSVP refresh reduction includes the following features:

- RSVP message bundling using the bundle message
- RSVP Message ID to reduce message processing overhead

- Reliable delivery of RSVP messages using Message ID, Message Ack, and Message Nack
- Summary refresh to reduce the amount of information transmitted every refresh interval

The RSVP refresh reduction specification (RFC 2961) allows you to enable some or all of the above capabilities on a router. It also describes various procedures that a router can use to detect the refresh reduction capabilities of its neighbor.

The Junos OS supports all of the refresh reduction extensions, some of which can be selectively enabled or disabled. The Junos OS supports Message ID and therefore can perform reliable message delivery only for Path and Resv messages.

For information about how to configure RSVP refresh reduction, see [“Configuring RSVP Refresh Reduction” on page 473](#).

---

## MTU Signaling in RSVP

The maximum transmission unit (MTU) is the largest size packet or frame, in bytes, that can be sent in a network. An MTU that is too large might cause retransmissions. Too small an MTU might cause the router to send and handle relatively more header overhead and acknowledgments. There are default values for MTUs associated with various protocols. You can also explicitly configure an MTU on an interface.

When an LSP is created across a set of links with different MTU sizes, the ingress router does not know what the smallest MTU is on the LSP path. By default, the MTU for an LSP is 1,500 bytes.

If this MTU is larger than the MTU of one of the intermediate links, traffic might be dropped, because MPLS packets cannot be fragmented. Also, the ingress router is not aware of this type of traffic loss, because the control plane for the LSP would still function normally.

To prevent this type of packet loss in MPLS LSPs, you can configure MTU signaling in RSVP. This feature is described in RFC 3209. Juniper Networks supports the Integrated Services object for MTU signaling in RSVP. The Integrated Services object is described in RFCs 2210 and 2215. MTU signaling in RSVP is disabled by default.

To avoid packet loss due to MTU mismatches, the ingress router needs to do the following:

- Signal the MTU on the RSVP LSP—To prevent packet loss from an MTU mismatch, the ingress router needs to know what the smallest MTU value is along the path taken by the LSP. Once this MTU value is obtained, the ingress router can assign it to the LSP.
- Fragment packets—Using the assigned MTU value, packets that exceed the size of the MTU can be fragmented into smaller packets on the ingress router before they are encapsulated in MPLS and sent over the RSVP-signaled LSP.

Once both MTU signaling and packet fragmentation have been enabled on an ingress router, any route resolving to an RSVP LSP on this router uses the signaled MTU value. For information about how to configure this feature, see [“Configuring MTU Signaling in RSVP” on page 484](#).

The following sections describe how MTU signaling in RSVP works:

- [How the Correct MTU Is Signaled in RSVP on page 469](#)
- [Determining an Outgoing MTU Value on page 469](#)
- [MTU Signaling in RSVP Limitations on page 470](#)

## How the Correct MTU Is Signaled in RSVP

How the correct MTU is signaled in RSVP varies depending on whether the network devices (for example, routers) explicitly support MTU signaling in RSVP or not.

If the network devices support MTU signaling in RSVP, the following occur when you enable MTU signaling:

- The MTU is signaled from the ingress router to the egress router by means of the Adspec object. Before forwarding this object, the ingress router enters the MTU value associated with the interface over which the path message is sent. At each hop in the path, the MTU value in the Adspec object is updated to the minimum of the received value and the value of the outgoing interface.
- The ingress router uses the traffic specification (Tspec) object to specify the parameters for the traffic it is going to send. The MTU value signaled for the Tspec object at the ingress router is the maximum MTU value (9192 bytes for M Series and T Series routers, 9500 bytes for PTX Series Packet Transport Routers). This value does not change en route to the egress router.
- When the Adspec object arrives at the egress router, the MTU value is correct for the path (meaning it is the smallest MTU value discovered). The egress router compares the MTU value in the Adspec object to the MTU value in the Tspec object. It signals the smaller MTU using the Flowspec object in the Resv message.
- When the Resv object arrives at the ingress router, the MTU value in this object is used as the MTU for the next hops that use the LSP.

In a network where there are devices that do not support MTU signaling in RSVP, you might have the following behaviors:

- If the egress router does not support MTU signaling in RSVP, the MTU is set to 1,500 bytes by default.
- A Juniper Networks transit router that does not support MTU signaling in RSVP sets an MTU value of 1,500 bytes in the Adspec object by default.

### Related Documentation

- [Configuring the Media MTU](#)

## Determining an Outgoing MTU Value

The outgoing MTU value is the smaller of the values received in the Adspec object compared to the MTU value of the outgoing interface. The MTU value of the outgoing interface is determined as follows:

- If you configure an MTU value under the **[family mpls]** hierarchy level, this value is signaled.
- If you do not configure an MTU, the **inet** MTU is signaled.

## MTU Signaling in RSVP Limitations

---

The following are limitations to MTU signaling in RSVP:

- Changes in the MTU value might cause a temporary loss of traffic in the following situations:
  - For link protection and node protection, the MTU of the bypass is only signaled at the time the bypass becomes active. During the time it takes for the new path MTU to be propagated, packet loss might occur because of an MTU mismatch.
  - For fast reroute, the MTU of the path is updated only after the detour becomes active, causing a delay in an update to the MTU at the ingress router. Until the MTU is updated, packet loss might occur if there is an MTU mismatch.

In both cases, only packets that are larger than the detour or bypass MTU are lost.

- When an MTU is updated, it triggers a change in the next hop. Any change in the next hop causes the route statistics to be lost.
- The minimum MTU supported for MTU signaling in RSVP is 1,488 bytes. This value prevents a false or incorrectly configured value from being used.
- For single-hop LSPs, the MTU value displayed by the **show** commands is the RSVP-signaled value. However, this MPLS value is ignored and the correct IP value is used.

## CHAPTER 10

# Configuring RSVP

- [Minimum RSVP Configuration on page 471](#)
- [Configuring RSVP and MPLS on page 472](#)
- [Configuring RSVP Interfaces on page 473](#)
- [Configuring RSVP Node ID Hellos on page 478](#)
- [Configuring Hello Acknowledgments for Nonsession RSVP Neighbors on page 479](#)
- [Switching LSPs Away from a Network Node on page 479](#)
- [Configuring RSVP Setup Protection on page 480](#)
- [Configuring Load Balancing Across RSVP LSPs on page 481](#)
- [Configuring RSVP Automatic Mesh on page 482](#)
- [Configuring Timers for RSVP Refresh Messages on page 483](#)
- [Preempting RSVP Sessions on page 484](#)
- [Configuring MTU Signaling in RSVP on page 484](#)
- [Configuring Ultimate-Hop Popping for LSPs on page 486](#)
- [Configuring RSVP to Pop the Label on the Ultimate-Hop Router on page 489](#)
- [Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs on page 490](#)
- [Tracing RSVP Protocol Traffic on page 491](#)

## Minimum RSVP Configuration

---

To enable RSVP on a single interface, include the **rsvp** statement and specify the interface using the **interface** statement. This is the minimum RSVP configuration. All other RSVP configuration statements are optional.

```
rsvp {
 interface interface-name;
}
```

You can include these statements at the following hierarchy levels:

- **[edit protocols]**
- **[edit logical-systems *logical-system-name* protocols]**

To enable RSVP on all interfaces, substitute **all** for the *interface-name* variable.

If you have configured interface properties on a group of interfaces and want to disable RSVP on one of the interfaces, include the **disable** statement:

```
interface interface-name {
 disable;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface interface-name** ]
- [edit logical-systems *logical-system-name* protocols rsvp **interface interface-name** ]

---

## Configuring RSVP and MPLS

The primary purpose of the Junos RSVP software is to support dynamic signaling within label-switched paths (LSPs). When you enable both MPLS and RSVP on a router, MPLS becomes a client of RSVP. No additional configuration is required to bind MPLS and RSVP.

You can configure MPLS to set up signaled paths by using the **label-switched-path** statement at the [edit protocols mpls] hierarchy level. Each LSP translates into a request for RSVP to initiate an RSVP session. This request is passed through the internal interface between label switching and RSVP. After examining the request information, checking RSVP states, and checking the local routing tables, RSVP initiates one session for each LSP. The session is sourced from the local router and is destined for the target of the LSP.

When an RSVP session is successfully created, the LSP is set up along the paths created by the RSVP session. If the RSVP session is unsuccessful, RSVP notifies MPLS of its status. It is up to MPLS to initiate backup paths or continue retrying the initial path.

To pass label-switching signaling information, RSVP supports four additional objects: Label Request object, Label object, Explicit Route object, and Record Route object. For an LSP to be set up successfully, all routers along the path must support MPLS, RSVP, and the four objects. Of the four objects, the Record Route object is not mandatory.

To configure MPLS and make it a client of RSVP, do the following:

- Enable MPLS on all routers that will participate in the label switching (this is, on all routers that might be part of a label-switching path).
- Enable RSVP on all routers and on all router interfaces that form the LSP.
- Configure the routers at the beginning of the LSP.

### Example: Configuring RSVP and MPLS

The following shows a sample configuration for a router at the beginning of an LSP:

```
[edit]
protocols {
 mpls {
 label-switched-path sf-to-london {
```

```

 to 192.168.1.4;
 }
}
rsvp {
 interface so-0/0/0;
}
}

```

The following shows a sample configuration for all the other routers that form the LSP:

```

[edit]
protocols {
 mpls {
 interface so-0/0/0;
 }
 rsvp {
 interface so-0/0/0;
 }
}

```

## Configuring RSVP Interfaces

The following sections describe how to configure RSVP interfaces:

- [Configuring RSVP Refresh Reduction on page 473](#)
- [Configuring the RSVP Hello Interval on page 475](#)
- [Configuring RSVP Authentication on page 476](#)
- [Configuring the Bandwidth Subscription for Class Types on page 476](#)
- [Configuring the RSVP Update Threshold on an Interface on page 476](#)
- [Configuring RSVP for Unnumbered Interfaces on page 477](#)

### Configuring RSVP Refresh Reduction

You can configure RSVP refresh reduction on each interface by including the following statements in the interface configuration:

- **aggregate** and **reliable**—Enable all RSVP refresh reduction features: RSVP message bundling, RSVP message ID, reliable message delivery, and summary refresh.  
In order to have refresh reduction and reliable delivery, you must include the **aggregate** and **reliable** statements.
- **no-aggregate**—Disable RSVP message bundling and summary refresh.
- **no-reliable**—Disable RSVP message ID, reliable message delivery, and summary refresh.

For more information on RSVP refresh reduction, see [“RSVP Refresh Reduction” on page 467](#).

If the **no-reliable** statement is configured on the router (reliable message delivery is disabled), the router accepts RSVP messages that include the Message ID object but ignores the Message ID object and continues performing standard message processing. No error is generated in this case, and RSVP operates normally.

However, not all combinations between two neighbors with different refresh reduction capabilities function correctly. For example, a router is configured with either the **aggregate** statement and **no-reliable** statement or with the **reliable** and **no-aggregate** statements. If an RSVP neighbor sends a Summary Refresh object to this router, no error is generated, but the Summary Refresh object cannot be processed. Consequently, RSVP states can time out on this router if the neighbor is relying only on Summary Refresh to refresh those RSVP states.

We recommend, unless there are specific requirements, that you configure RSVP refresh reduction in a similar manner on each RSVP neighbor.

To enable all RSVP refresh reduction features on an interface, include the **aggregate** statement:

```
aggregate;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name*]

To disable RSVP message bundling and summary refresh, include the **no-aggregate** statement:

```
no-aggregate;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name*]

To enable RSVP message ID and reliable message delivery on an interface, include the **reliable** statement:

```
reliable;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name*]

To disable RSVP message ID, reliable message delivery, and summary refresh, include the **no-reliable** statement:

```
no-reliable;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name*]



### Determining the Refresh Reduction Capability of RSVP Neighbors

To determine the RSVP refresh reduction capability of an RSVP neighbor, you need the following information:

- The RR bit advertised by the neighbor
- The local configuration of RSVP refresh reduction
- The actual RSVP messages received from the neighbor

To obtain this information, you can issue a **show rsvp neighbor detail** command. Sample output follows:

```
user@host> show rsvp neighbor detail
RSVP neighbor: 6 learned
 Address: 192.168.224.178 via: fxp1.0 status: Up
 Last changed time: 10:06, Idle: 5 sec, Up cnt: 1, Down cnt: 0
 Message received: 36
 Hello: sent 69, received: 69, interval: 9 sec
 Remote instance: 0x60b8feba, Local instance: 0x74bc7a8d
 Refresh reduction: not operational

 Address: 192.168.224.186 via: fxp2.0 status: Down
 Last changed time: 10:17, Idle: 40 sec, Up cnt: 0, Down cnt: 0
 Message received: 6
 Hello: sent 20, received: 0, interval: 9 sec
 Remote instance: 0x0, Local instance: 0x2ae1b339
 Refresh reduction: incomplete
 Remote end: disabled, Ack-extension: enabled

 Address: 192.168.224.188 via: fxp2.0 status: Up
 Last changed time: 4:15, Idle: 0 sec, Up cnt: 1, Down cnt: 0
 Message received: 55
 Hello: sent 47, received: 31, interval: 9 sec
 Remote instance: 0x6436a35b, Local instance: 0x663849f0
 Refresh reduction: operational
 Remote end: enabled, Ack-extension: enabled
```

For more information on the **show rsvp neighbor detail** command.

### Configuring the RSVP Hello Interval

RSVP monitors the status of the interior gateway protocol (IGP) (IS-IS or OSPF) neighbors and relies on the IGP protocols to detect when a node fails. If an IGP protocol declares a neighbor down (because hello packets are no longer being received), RSVP also brings down that neighbor. However, the IGP protocols and RSVP still act independently when bringing a neighbor up.

For Juniper Networks routers, configuring a shorter or longer RSVP hello interval has no impact on whether or not an RSVP session is brought down. RSVP sessions are kept up even if RSVP hello packets are no longer being received. RSVP sessions are maintained until either the router stops receiving IGP hello packets or the RSVP Path and Resv messages time out.

However, the RSVP hello interval might impact when another vendor's equipment brings down an RSVP session. For example, a neighboring non-Juniper Networks router might be configured to monitor RSVP hello packets.

To modify how often RSVP sends hello packets, include the **hello-interval** statement:

```
hello-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

## Configuring RSVP Authentication

All RSVP protocol exchanges can be authenticated to guarantee that only trusted neighbors participate in setting up reservations. By default, RSVP authentication is disabled.

RSVP authentication uses a Hashed Message Authentication Code (HMAC)-MD5 message-based digest. This scheme produces a message digest based on a secret authentication key and the message contents. (The message contents also include a sequence number.) The computed digest is transmitted with RSVP messages. Once you have configured authentication, all received and transmitted RSVP messages with all neighbors are authenticated on this interface.

MD5 authentication provides protection against forgery and message modification. It also can prevent replay attacks. However, it does not provide confidentiality, because all messages are sent in clear text.

By default, authentication is disabled. To enable authentication, configure a key on each interface by including the **authentication-key** statement:

```
authentication-key key;
```

You can include this statement at the following hierarchy levels:

- [edit protocols **rsvp** **interface** *interface-name*]
- [edit logical-systems *logical-system-name* protocols **rsvp** **interface** *interface-name*]

## Configuring the Bandwidth Subscription for Class Types

By default, RSVP allows 100 percent of the bandwidth for a class type to be used for RSVP reservations. When you oversubscribe a class type for a multiclass LSP, the aggregate demand of all RSVP sessions is allowed to exceed the actual capacity of the class type.

For detailed instructions on how to configure the bandwidth subscription for class types, see [“Configuring the Bandwidth Subscription Percentage for LSPs” on page 323](#).

## Configuring the RSVP Update Threshold on an Interface

The interior gateway protocols (IGPs) maintain the traffic engineering database, but the current available bandwidth on the traffic engineering database links originates from RSVP. When a link's bandwidth changes, RSVP informs the IGPs, which can then update the traffic engineering database and forward the new bandwidth information to all

network nodes. The network nodes then know how much bandwidth is available on the traffic engineering database link (local or remote), and CSPF can correctly compute the paths.

However, IGP updates can consume excessive system resources. Depending on the number of nodes in a network, it might not be desirable to perform an IGP update for small changes in bandwidth. By configuring the **update-threshold** statement at the **[edit protocols rsvp]** hierarchy level, you can adjust the threshold at which a change in the reserved bandwidth triggers an IGP update.

You can configure a value of from 1 percent through 20 percent (the default is 10 percent) for when to trigger an IGP update. If the change in the reserved bandwidth is greater than or equal to the configured threshold percentage of the static bandwidth on that interface, then an IGP update occurs. For example, if you have configured the **update-threshold** statement to be 15 percent and the router discovers that the reserved bandwidth on a link has changed by 10 percent of the link bandwidth, RSVP does not trigger an IGP update. However, if the reserved bandwidth on a link changes by 20 percent of the link bandwidth, RSVP triggers an IGP update.

To adjust the threshold at which a change in the reserved bandwidth triggers an IGP update, include the **update-threshold** statement:

```
update-threshold percentage;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols rsvp interface interface-name]**
- **[edit logical-systems logical-system-name protocols rsvp interface interface-name]**

Because of the update threshold, it is possible for Constrained Shortest Path First (CSPF) to compute a path using outdated traffic engineering database bandwidth information on a link. If RSVP attempts to establish an LSP over that path, it might find that there is insufficient bandwidth on that link. When this happens, RSVP triggers an IGP traffic engineering database update, flooding the updated bandwidth information on the network. CSPF can then recompute the path by using the updated bandwidth information, and attempt to find a different path, avoiding the congested link. Note that this functionality is the default and does not need any additional configuration.

You can configure the **rsvp-error-hold-time** statement at the **[edit protocols mpls]** hierarchy level or the **[edit logical-systems logical-system-name protocols mpls]** hierarchy level to improve the accuracy of the traffic engineering database (including the accuracy of bandwidth estimates for LSPs) using information provided by PathErr messages. See [“Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages” on page 68](#).

## Configuring RSVP for Unnumbered Interfaces

The Junos OS supports RSVP traffic engineering over unnumbered interfaces. Traffic engineering information about unnumbered links is carried in the IGP traffic engineering extensions for OSPF and IS-IS as described in RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*, and RFC 4205, *Intermediate System to Intermediate System (IS-IS) Extensions in Support of Generalized Multi-Protocol Label*

*Switching (GMPLS)*. Unnumbered links can also be specified in the MPLS traffic engineering signaling as described in RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*. This feature allows you avoid having to configure IP addresses for each interface participating in the RSVP-signaled network.

To configure RSVP for unnumbered interfaces, you must configure the router with a router ID using the **router-id** statement specified at the **[edit routing-options]** hierarchy level. The router ID must be available for routing (you can typically use the loopback address). The RSVP control messages for the unnumbered links are sent using the router ID address (rather than a randomly selected address).

To configure link protection and fast reroute on a router with unnumbered interfaces enabled, you must configure at least two addresses. We recommend that you configure a secondary interface on the loopback in addition to configuring the router ID.

---

## Configuring RSVP Node ID Hellos

You can configure node-ID based RSVP hellos to ensure that Juniper Networks routers can interoperate with the equipment of other vendors. By default, the Junos OS uses interface-based RSVP hellos. Node-ID based RSVP hellos are specified in RFC 4558, *Node-ID Based Resource Reservation Protocol (RSVP) Hello: A Clarification Statement*. RSVP node-ID hellos are useful if you have configured BFD to detect problems over RSVP interfaces, allowing you to disable interface-hellos for these interfaces. You can also use node-ID hellos for graceful-restart procedures.

Node-ID hellos can be enabled globally for all RSVP neighbors. By default, node-ID hello support is disabled. If you have not enabled RSVP node IDs on the router, the Junos OS does not accept any node-ID hello packets.

To enable RSVP node-ID hellos globally on the router, include the **node-hello** statement:

```
node-hello;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols rsvp]**
- **[edit logical-systems *logical-systems-name* protocols rsvp]**

You can also explicitly disable RSVP interface hellos globally. This type of configuration might be necessary in networks where the Juniper Networks router has numerous RSVP connections with equipment from other vendors. However, if you disable RSVP interface hellos globally, you can also configure a hello interval on an RSVP interface using the **hello-interval** statement. This configuration disables RSVP interface hellos globally, but enables RSVP interface hellos on the specified interface (the RSVP interface you configure the **hello-interval** statement on). This configuration might be necessary in a heterogeneous network in which some devices support RSVP node ID hellos and other devices support RSVP interface hellos.

To disable RSVP interface hellos globally on the router, include the **no-interface-hello** statement:

```
no-interface-hello;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]
- [edit logical-systems *logical-systems-name* protocols rsvp]

## Configuring Hello Acknowledgments for Nonsession RSVP Neighbors

The **hello-acknowledgements** statement controls the hello acknowledgment behavior between RSVP neighbors regardless of whether or not they are in the same session.

Hello messages received from RSVP neighbors that are not part of a common RSVP session are discarded. If you configure the **hello-acknowledgements** statement at the [edit protocols rsvp] hierarchy level, hello messages from nonsession neighbors are acknowledged with a hello acknowledgment message. When hellos are received from nonsession neighbors, an RSVP neighbor relationship is created and periodic hello messages can now be received from the nonsession neighbor. The **hello-acknowledgements** statement is disabled by default. Configuring this statement allows RSVP-capable routers to be discovered using hello packets and verifies that the interface is able to receive RSVP packets before sending any MPLS LSP setup messages.

Once you enable hello acknowledgments for nonsession RSVP neighbors, the router continues to acknowledge hello messages from any nonsession RSVP neighbors unless the interface itself goes down or you change the configuration. Interface-based neighbors are not automatically aged out.

```
hello-acknowledgements;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]
- [edit logical-systems *logical-system-name* protocols rsvp]

## Switching LSPs Away from a Network Node

You can configure the router to switch active LSPs away from a network node using a bypass LSP enabled for an interface. This feature might be used to maintain active networks when a device needs to be replaced without interrupting traffic transiting the network. The LSPs can be either static or dynamic.

1. You first need to configure either link or node protection for the traffic that needs to pass around the network device you intend to disable. To function properly, the bypass LSP must use a different logical interface than the protected LSP.
2. To prepare the router to begin switching traffic away from a network node, configure the **always-mark-connection-protection-tlv** statement:

```
always-mark-connection-protection-tlv;
```

The router then marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality.

You can configure this statement at the following hierarchy levels:

- **[edit protocols mpls interface *interface-name*]**
  - **[edit logical-systems *logical-system-name* protocols mpls interface *interface-name*]**
3. You then need to configure the **switch-away-lsps** statement to switch the traffic from the protected LSP to the bypass LSP, effectively bypassing the default downstream network device. The actual link itself is not brought down by this configuration.

To configure the router to switch traffic away from a network node, configure the **switch-away-lsps** statement:

```
switch-away-lsps;
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols mpls interface *interface-name*]**
- **[edit logical-systems *logical-system-name* protocols mpls interface *interface-name*]**

Note the following limitations related to switching active LSPs away from a network node:

- The switch-away feature is supported on MX Series routers only.
- The switch-away feature is not supported for switching traffic from primary point-to-multipoint LSPs to bypass point-to-multipoint LSPs. If you configure the **switch-away-lsps** statement for a point-to-multipoint LSP, traffic is not switched to the bypass point-to-multipoint LSP.
- If you configure the switch-away feature on an interface along the path of a dynamic LSP, new dynamic LSPs cannot be established over that path. The switch-away feature prevents the make-before-break behavior of RSVP-signaled LSPs. The make-before-break behavior normally causes the router to first attempt to re-signal a dynamic LSP before tearing down the original.

**Related  
Documentation**

- [Configuring Node Protection or Link Protection for LSPs on page 509](#)

---

## Configuring RSVP Setup Protection

You can configure the facility-backup fast reroute mechanism to provide setup protection for LSPs which are in the process of being signaled. Both point-to-point LSPs and point-to-multipoint LSPs are supported. This feature is applicable in the following scenario:

1. A failed link or node is present on the strict explicit path of an LSP before the LSP is signaled.
2. There is also a bypass LSP protecting the link or node.

3. RSVP signals the LSP through the bypass LSP. The LSP appears as if it was originally set up along its primary path and then failed over to the bypass LSP because of the link or node failure.
4. When the link or node has recovered, the LSP can be automatically reverted to the primary path.

You should configure the **setup-protection** statement at the **[edit protocols rsvp]** on each of the routers along the LSP path on which you want to enable LSP setup protection. You should also configure IGP traffic engineering on all of the routers on the LSP path. You can issue a **show rsvp session** command to determine whether or not the LSP has setup protection enabled on a router acting as a point of local repair (PLR) or a merge point.

To enable RSVP setup protection, include the **setup-protection** statement

```
setup-protection;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols rsvp]**
- **[edit logical-systems *logical-system-name* protocols rsvp]**

## Configuring Load Balancing Across RSVP LSPs

By default, when you have configured several RSVP LSPs to the same egress router, the LSP with the lowest metric is selected and carries all traffic. If all of the LSPs have the same metric, one of the LSPs is selected at random and all traffic is forwarded over it.

Alternatively, you can load-balance traffic across all of the LSPs by enabling per-packet load balancing.

To enable per-packet load balancing on an ingress LSP, configure the **policy-statement** statement as follows:

```
[edit policy-options]
policy-statement policy-name {
 then {
 load-balance per-packet;
 }
 accept;
}
```

You then need to apply this statement as an export policy to the forwarding table.

Once per-packet load balancing is applied, traffic is distributed equally between the LSPs (by default).

You need to configure per-packet load balancing if you want to enable PFE fast reroute. To enable PFE fast reroute, include the **policy-statement** statement for per-packet load balancing shown in this section in the configuration of each of the routers where a reroute might take place. See also [“Configuring Fast Reroute” on page 226](#).

You can also load-balance the traffic between the LSPs in proportion to the amount of bandwidth configured for each LSP. This capability can better distribute traffic in networks with asymmetric bandwidth capabilities across external links, since the configured bandwidth of an LSP typically reflects the traffic capacity of that LSP.

To configure RSVP LSP load balancing, include the **load-balance** statement with the **bandwidth** option:

```
load-balance {
 bandwidth;
}
```

You can configure this statement at the following hierarchy levels:

- [edit protocols **rsvp**]
- [edit logical-systems *logical-system-name* protocols **rsvp**]

Keep the following information in mind when you use the **load-balance** statement:

- If you configure the **load-balance** statement, the behavior of currently running LSPs is not altered. To force currently running LSPs to use the new behavior, you can issue a **clear mpls lsp** command.
- The **load-balance** statement only applies to ingress LSPs that have per-packet load balancing enabled.
- For Differentiated Services-aware traffic engineered LSPs, the bandwidth of an LSP is calculated by summing the bandwidth of all of the class types.

---

## Configuring RSVP Automatic Mesh

---

You can configure RSVP to establish point-to-point label-switched paths (LSPs) automatically for any new PE router added to a full mesh of LSPs. To enable this feature, you must configure the **rsvp-te** statement on all of the PE routers in the full mesh.



**NOTE:** You cannot configure RSVP automatic mesh in conjunction with CCC. CCC cannot use the dynamically generated LSPs.

To configure RSVP automatic mesh, include the **rsvp-te** statement:

```
rsvp-te {
 destination-networks network-prefix;
 label-switched-path-template (Multicast) {
 default-template;
 template-name;
 }
}
```

You can configure these statements at the following hierarchy levels:

- [edit routing-options dynamic-tunnels *tunnel-name*]
- [edit logical-systems *logical-system-name* routing-options dynamic-tunnels *tunnel-name*]



You must also configure the following statements to enable RSVP automatic mesh:

- **destination-networks**—Specify the IP version 4 (IPv4) prefix range for the destination network. Dynamic tunnels within the specified IPv4 prefix range can be initiated.
- **label-switched-path-template (Multicast)**—You can configure either the default template explicitly using the **default-template** option, or you can configure an LSP template of your own using the **template-name** option. The LSP template acts as a model configuration for the dynamically generated LSPs.

**Related  
Documentation**

- *Flooding Unknown Traffic Using Point-to-Multipoint LSPs in VPLS*
- *Example: Configuring RSVP Automatic Mesh*

## Configuring Timers for RSVP Refresh Messages

RSVP uses two related timing parameters:

- **refresh-time**—The refresh time controls the interval between the generation of successive refresh messages. The default value for the refresh time is 45 seconds. This number is derived from the **refresh-time** statement's default value of 30, multiplied by a fixed value of 1.5. This computation differs from RFC 2205, which states that the refresh time should be multiplied by a random value in the range from 0.5 through 1.5.

Refresh messages include path and Resv messages. Refresh messages are sent periodically so that reservation states in neighboring nodes do not time out. Each path and Resv message carries the refresh timer value, and the receiving node extracts this value from the messages.

- **keep-multiplier**—The keep multiplier is a small, locally configured integer from 1 through 255. The default value is 3. It indicates the number of messages that can be lost before a particular state is declared stale and must be deleted. The keep multiplier directly affects the lifetime of an RSVP state.

To determine the lifetime of a reservation state, use the following formula:

$$\text{lifetime} = (\text{keep-multiplier} + 0.5) \times (1.5 \times \text{refresh-time})$$

In the worst case, (**keep-multiplier** – 1) successive refresh messages must be lost before a reservation state is deleted.

We do not recommend configuring a short RSVP hello timer. If quick discovery of a failed neighbor is needed, configure short IGP (OSPF or IS-IS) hello timers.

By default, the refresh timer value is 30 seconds. To modify this value, include the **refresh-time** statement:

```
refresh-time seconds;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols rsvp]**
- **[edit logical-systems *logical-system-name* protocols rsvp]**

The default value of the keep multiplier is 3. To modify this value, include the **keep-multiplier** statement:

```
keep-multiplier number;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]
- [edit logical-systems *logical-system-name* protocols rsvp]

---

## Preempting RSVP Sessions

Whenever bandwidth is insufficient to handle all RSVP sessions, you can control the preemption of RSVP sessions. By default, an RSVP session is preempted only by a new higher-priority session.

To always preempt a session when the bandwidth is insufficient, include the **preemption** statement with the **aggressive** option:

```
preemption aggressive;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]
- [edit logical-systems *logical-system-name* protocols rsvp]

To disable RSVP session preemption, include the **preemption** statement with the **disabled** option:

```
preemption disabled;
```

To return to the default (that is, preempt a session only for a new higher-priority session), include the **preemption** statement with the **normal** option:

```
preemption normal;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp]
- [edit logical-systems *logical-system-name* protocols rsvp]

---

## Configuring MTU Signaling in RSVP

To configure maximum transmission unit (MTU) signaling in RSVP, you need to configure MPLS to allow IP packets to be fragmented before they are encapsulated in MPLS. You also need to configure MTU signaling in RSVP. For troubleshooting purposes, you can configure MTU signaling alone without enabling packet fragmentation.

To configure MTU signaling in RSVP, include the **path-mtu** statement:

```
path-mtu {
 allow-fragmentation;
 rsvp {
```

```

 mtu-signaling;
 }
}

```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

The following sections describe how to enable packet fragmentation and MTU signaling in RSVP:

- [Enabling MTU Signaling in RSVP on page 485](#)
- [Enabling Packet Fragmentation on page 485](#)

## Enabling MTU Signaling in RSVP

To enable MTU signaling in RSVP, include the **rsvp mtu-signaling** statement:

```
rsvp mtu-signaling;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls path-mtu]
- [edit logical-systems *logical-system-name* protocols mpls path-mtu]

Once you have committed the configuration, changes in the MTU signaling behavior for RSVP take effect the next time the path is refreshed.

You can configure the **mtu-signaling** statement by itself at the [edit protocols mpls path-mtu rsvp] hierarchy level. This can be useful for troubleshooting. If you configure just the **mtu-signaling** statement, you can use the **show rsvp session detail** command to determine what the smallest MTU is on an LSP. The **show rsvp session detail** command displays the MTU value received and sent in the Adspec object.

## Enabling Packet Fragmentation

To allow IP packets to be fragmented before they are encapsulated in MPLS, include the **allow-fragmentation** statement:

```
allow-fragmentation;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls path-mtu]
- [edit logical-systems *logical-system-name* protocols mpls path-mtu]




**NOTE:** Do not configure the **allow-fragmentation** statement alone. Always configure it in conjunction with the **mtu-signaling** statement.

## Configuring Ultimate-Hop Popping for LSPs

By default, RSVP-signaled LSPs use penultimate-hop popping (PHP).


Figure 28 on page 222 illustrates a penultimate-hop popping LSP between Router PE1 and Router PE2. Router CE1 forwards a packet to its next hop (Router PE1), which is also the LSP ingress. Router PE1 pushes label 1 on the packet and forwards the labeled packet to Router P1. Router P1 completes the standard MPLS label swapping operation, swapping label 1 for label 2, and forwards the packet to Router P2. Since Router P2 is the penultimate-hop router for the LSP to Router PE2, it first pops the label and then forwards the packet to Router PE2. When Router PE2 receives it, the packet can have a service label, an explicit-null label, or just be a plain IP or VPLS packet. Router PE2 forwards the unlabeled packet to Router CE2.

Figure 40: Penultimate-Hop Popping for an LSP



You can also configure ultimate-hop popping (UHP) (as shown in Figure 29 on page 222) for RSVP-signaled LSPs. Some network applications can require that packets arrive at the egress router (Router PE2) with a non-null outer label. For an ultimate-hop popping LSP, the penultimate router (Router P2 in Figure 29 on page 222) performs the standard MPLS label swapping operation (in this example, label 2 for label 3) before forwarding the packet to egress Router PE2. Router PE2 pops the outer label and performs a second lookup of the packet address to determine the end destination. It then forwards the packet to the appropriate destination (either Router CE2 or Router CE4).

Figure 41: Ultimate-Hop Popping for an LSP



The following network applications require that you configure UHP LSPs:

- MPLS-TP for performance monitoring and in-band OAM
- Edge protection virtual circuits

The following features do not support the UHP behavior:

- LDP-signaled LSPs
- Static LSPs

- Point-to-multipoint LSPs
- CCC
- **traceroute** command

For more information about UHP behavior, see Internet draft draft-ietf-mpls-rsvp-te-no-php-oob-mapping-01.txt, *Non PHP behavior and Out-of-Band Mapping for RSVP-TE LSPs*.

For point-to-point RSVP-signaled LSPs, UHP behavior is signaled from the LSP ingress. Based on the ingress router configuration, RSVP can signal the UHP LSP with the non-PHP flag set. RSVP PATH messages carry the two flags in the LSP-ATTRIBUTES object. When the egress router receives the PATH message, it assigns a non-null label to the LSP. RSVP also creates and installs two routes in the mpls.0 routing table. S refers to the S bit of the MPLS label, which indicates whether or not the bottom of the label stack has been reached.

- Route S=0—Indicates that there are more labels in the stack. The next hop for this route points to the mpls.0 routing table, triggering a chained MPLS label lookup to discover the remaining MPLS labels in the stack.
- Route S=1—Indicates that there are no more labels. The next hop points to the inet.0 routing table if the platform supports chained and multi-family lookup. Alternatively, the label route can point to a VT interface to initiate IP forwarding.

If you enable UHP LSPs, MPLS applications such as Layer 3 VPNs, VPLS, Layer 2 VPNs, and Layer 2 circuits can use the UHP LSPs. The following explains how UHP LSPs affect the different types of MPLS applications:

- Layer 2 VPNs and Layer 2 circuits—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label (S=0) is the UHP label, and the inner label (S=1) is the VC label. A lookup based on the transport label results in a table handle for the mpls.0 routing table. There is an additional route in the mpls.0 routing table corresponding to the inner label. A lookup based on the inner label results in the CE router next hop.
- Layer 3 VPN—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label (S=0) is the UHP label, and the inner label is the VPN label (S=1). A lookup based on the transport label results in the table handle for the mpls.0 routing table. There are two cases in this scenario. By default, Layer 3 VPNs advertise the per-next hop label. A lookup based on the inner label results in the next hop toward the CE router. However, if you have configured the **vrf-table-label** statement for the Layer 3 VPN routing instance, the inner LSI label points to the VRF routing table. An IP lookup is also completed for the VRF routing table.



**NOTE:** UHP for Layer 3 VPNs configured with the **vrf-table-label** statement is supported on MX 3D Universal Edge Routers only.

- VPLS—A packet arrives at the PE router (egress of the UHP LSP) with two labels. The outer label is the transport label (S=0) and the inner label is the VPLS label (S=1). A

lookup based on the transport label results in the table handle for the mpls.0 routing table. A lookup based on the inner label in mpls.0 routing table results in the LSI tunnel interface of the VPLS routing instance if tunnel-services is not configured (or a VT interface not available). MX 3D Series routers support chained lookup and multi-family lookup.



**NOTE:** UHP for VPLS configured with the `no-tunnel-service` statement is supported on MX 3D Series routers only.

- **IPv4 over MPLS**—A packet arrives at the PE router (egress of the UHP LSP) with one label (S=1). A lookup based on this label returns a VT tunnel interface. Another IP lookup is completed on the VT interface to determine where to forward the packet. If the routing platform supports multi-family and chained lookups (for example, MX 3D routers and PTX Series Packet Transport Routers), lookup based on label route (S=1) points to the inet.0 routing table.
- **IPv6 over MPLS**—For IPv6 tunneling over MPLS, PE routers advertise IPv6 routes to each other with a label value of 2. This is the explicit null label for IPv6. As a result, the forwarding next hops for IPv6 routes that are learned from remote PE routers normally push two labels. The inner label is 2 (it could be different if the advertising PE router is from another vendor), and the router label is the LSP label. Packets arrive at the PE router (egress of the UHP LSP) with two labels. The outer label is the transport label (S=0), and the inner label is the IPv6 explicit-null label (label 2). Lookup based on the inner label in the mpls.0 routing table redirects back to the mpls.0 routing table. On MX 3D Series routers, the inner label (label 2) is stripped off and an IPv6 lookup is done using the inet6.0 routing table.
- **Enabling both PHP and UHP LSPs**—You can configure both PHP and UHP LSPs over the same network paths. You can separate PHP and UHP traffic by selecting forwarding LSP next hops using a regular expression with the `install-nexthop` statement. You can also separate traffic by simply naming the LSPs appropriately.

The following statements enable ultimate-hop popping for an LSP. You can enable this feature on a specific LSP or for all of the ingress LSPs configured on the router. Configure these statements on the router at the LSP ingress.

1. To enable ultimate-hop popping, include the **ultimate-hop-popping** statement:

**ultimate-hop-popping;**

Include this statement at the `[edit protocols mpls label-switched-path label-switched-path-name]` hierarchy level to enable ultimate-hop popping on a specific LSP. Include this statement at the `[edit protocols mpls]` hierarchy level to enable ultimate-hop popping on all of the ingress LSPs configured on the router. You can also configure the **ultimate-hop-popping** statement under the equivalent `[edit logical-routers]` hierarchy levels.



**NOTE:** When you enable ultimate-hop popping, RSVP attempts to resignal existing LSPs as ultimate-hop popping LSPs in a make-before-break fashion. If an egress router does not support ultimate-hop popping, the existing LSP is torn down (RSVP sends a PathTear message along an LSP's path, removing the path state and dependent reservation state and releasing the associated networking resources).

If you disable ultimate-hop popping, RSVP resignals existing LSPs as penultimate-hop popping LSPs in a make-before-break fashion.

2. If you want to enable both ultimate-hop-popping and chained next hops on MX 3D Series routers only, you also need to configure the **enhanced-ip** option for the **network-services** statement:

```
network-services enhanced-ip;
```

You configure this statement at the **[edit chassis]** hierarchy level. Once you have configured the **network-services** statement, you need to reboot the router to enable UHP behavior.

#### Related Documentation

- [MPLS Label Allocation on page 26](#)
- [Configuring Corouted Bidirectional LSPs on page 220](#)
- [network-services](#)
- [ultimate-hop-popping on page 965](#)

## Configuring RSVP to Pop the Label on the Ultimate-Hop Router

You can control the label value advertised on the egress router of an LSP. The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop router removes the label and sends the packet to the egress router. When ultimate-hop popping is enabled, label 0 (IP version 4 [IPv4] Explicit Null label) is advertised. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label.



**NOTE:** Juniper Networks routers queue packets based on the incoming label. Routers from other vendors might queue packets differently. Keep this in mind when working with networks containing routers from multiple vendors.

To configure ultimate-hop popping for RSVP, include the **explicit-null** statement:

```
explicit-null;
```

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls]**
- **[edit logical-systems *logical-system-name* protocols mpls]**

- Related Documentation**
- [MPLS Label Overview on page 24](#)
  - [MPLS Label Allocation on page 26](#)

---

## Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs

By default, for both point-to-point and point-to-multipoint LSPs, penultimate-hop popping is used for MPLS traffic. MPLS labels are removed from packets on the router just before the egress router of the LSP. The plain IP packets are then forwarded to the egress router. For ultimate-hop popping, the egress router is responsible for both removing the MPLS label and processing the plain IP packet.

It can be beneficial to enable ultimate-hop popping on point-to-multipoint LSPs, particularly when transit traffic is traversing the same egress device. If you enable ultimate-hop popping, a single copy of traffic can be sent over the incoming link, saving significant bandwidth. By default, ultimate-hop popping is disabled.

You enable ultimate-hop popping for point-to-multipoint LSPs by configuring the **tunnel-services** statement. When you enable ultimate-hop popping, the Junos OS selects one of the available virtual loopback tunnel (VT) interfaces to loop back the packets to the PFE for IP forwarding. By default, the VT interface selection process is performed automatically. Bandwidth admission control is used to limit the number of LSPs that can be used on one VT interface. Once all the bandwidth is consumed on one interface, the Junos OS selects another VT interface with sufficient bandwidth for admission control.

If an LSP requires more bandwidth than is available from any of the VT interfaces, ultimate-hop popping cannot be enabled and penultimate-hop popping is enabled instead.

For ultimate-hop popping on point-to-multipoint LSPs to function properly, the egress router must have a PIC that provides tunnel services, such as the tunnel services PIC or the adaptive services PIC. Tunnel services are needed for popping the final MPLS label and for returning packets for IP address lookups.

You can explicitly configure which VT interfaces handle the RSVP traffic by including the **devices** option for the **tunnel-services** statement. The **devices** option allows you to specify which VT interfaces are to be used by RSVP. If you do not configure this option, all of the VT interfaces available to the router can be used.

To enable ultimate-hop popping for the egress point-to-multipoint LSPs on a router, configure the **tunnel-services** statement:

```
tunnel-services {
 devices device-names;
}
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols rsvp]**
- **[edit logical-systems *logical-system-name* protocols rsvp]**



To enable ultimate-hop popping for egress point-to-multipoint LSPs, you must also configure the **interface** statement with the **all** option:

```
interface all;
```

You must configure this statement at the **[edit protocols rsvp]** hierarchy level.

## Tracing RSVP Protocol Traffic

To trace RSVP protocol traffic, include the **traceoptions** statement:

```
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols rsvp]**
- **[edit logical-systems *logical-system-name* protocols rsvp]**

You can specify the following RSVP-specific flags in the RSVP **traceoptions** statement:

Use the **file** statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory **/var/log**. We recommend that you place RSVP tracing output in the file **rsvp-log**.

- **all**—All tracing operations.
- **error**—All detected error conditions
- **event**—RSVP-related events (helps to trace events related to RSVP graceful restart)
- **lmp**—RSVP-Link Management Protocol (LMP) interactions
- **packets**—All RSVP packets
- **path**—All path messages
- **pathtear**—PathTear messages
- **resv**—Resv messages
- **resvtear**—ResvTear messages
- **route**—Routing information
- **state**—Session state transitions, including when RSVP-signaled LSPs come up and go down.

To view the log file generated when you enable RSVP traceoptions, issue the **show log file-name** command, where **file-name** is the file you specified using the **traceoptions** statement.

For general information about tracing and global tracing options, see the *Junos OS Routing Protocols Library for Routing Devices*.

## Examples: Tracing RSVP Protocol Traffic

Trace RSVP path messages in detail:

```
[edit]
protocols {
 rsvp {
 traceoptions {
 file rsvp size 10m files 5;
 flag path;
 }
 }
}
```

Trace all RSVP messages:

```
[edit]
protocols {
 rsvp {
 traceoptions {
 file rsvp size 10m files 5;
 flag packets;
 }
 }
}
```

Trace all RSVP error conditions:

```
[edit]
protocols {
 rsvp {
 traceoptions {
 file rsvp size 10m files 5;
 flag error;
 }
 }
}
```

Trace RSVP state transitions:

```
[edit]
protocols {
 rsvp {
 traceoptions {
 file rsvp-data;
 flag state;
 }
 }
}
```

## RSVP Log File Output

The following is sample output generated by issuing the **show log *file-name*** command on a router on which RSVP traceoptions have been enabled with the **state** flag configured. The RSVP-signaled LSP E-D is shown being torn down on Mar 11 14:04:36.707092. On Mar 11 14:05:30.101492, it is shown coming back up.

```

user@host> show log rsvp-data
Mar 11 13:58:51 trace_on: Tracing to "/var/log/E/rsvp-data" started
Mar 11 13:58:51.286413 rsvp_iflchange for vt ifl vt-1/2/0.69206016
Mar 11 13:58:51.286718 RSVP add interface vt-1/2/0.69206016, ifindex 101, ifaddr
(null), family 1, is_appl_vt 0, already known
Mar 11 13:58:51.286818 RSVP Peer vt-1/2/0.69206016 TE-link __rpd:vt-1/2/0.69206016
Up
Mar 11 13:58:51.286978 RSVP add interface vt-1/2/0.69206016, ifindex 101, ifaddr
(null), family 3, is_appl_vt 0, already known
Mar 11 13:58:51.287962 RSVP add interface lt-1/2/0.2, ifindex 113, ifaddr (null),
family 2, is_appl_vt 0, already known
Mar 11 13:58:51.288629 RSVP add interface lt-1/2/0.2, ifindex 113, ifaddr 10.0.0.2,
family 1, is_appl_vt 0, already known
Mar 11 13:58:51.288996 RSVP add interface lt-1/2/0.17, ifindex 114, ifaddr (null),
family 2, is_appl_vt 0, already known
Mar 11 13:58:51.289593 RSVP add interface lt-1/2/0.17, ifindex 114, ifaddr (null),
family 3, is_appl_vt 0, already known
Mar 11 13:58:51.289949 RSVP add interface lt-1/2/0.17, ifindex 114, ifaddr
10.0.0.17, family 1, is_appl_vt 0, already known
Mar 11 13:58:51.290049 RSVP Peer lt-1/2/0.17 TE-link __rpd:lt-1/2/0.17 Up
Mar 11 13:59:05.042034 RSVP new bypass Bypass->10.0.0.18 on interface lt-1/2/0.17
to 10.0.0.18 avoid 0.0.0.0
Mar 11 14:04:36.707092 LSP "E-D" is Down (Reason: Reservation state deleted)
Session: 192.168.0.4(port/tunnel ID 10321 Ext-ID 192.168.0.5) Proto 0
Sender: 192.168.0.5(port/lsp ID 1)
Mar 11 14:04:36.707661 RSVP delete resv state, session 192.168.0.4(port/tunnel
ID 10321 Ext-ID 192.168.0.5) Proto 0
Mar 11 14:04:36.781185 RSVP delete path state, session 192.168.0.4(port/tunnel
ID 10321 Ext-ID 192.168.0.5) Proto 0
Mar 11 14:04:36.781440 RSVP delete session 192.168.0.4(port/tunnel ID 10321 Ext-ID
192.168.0.5) Proto 0
Mar 11 14:05:30.101492 RSVP new Session 192.168.0.4(port/tunnel ID 10321 Ext-ID
192.168.0.5) Proto 0, session ID 3
Mar 11 14:05:30.101722 RSVP new path state, session 192.168.0.4(port/tunnel ID
10321 Ext-ID 192.168.0.5) Proto 0
Mar 11 14:05:30.179124 RSVP new resv state, session 192.168.0.4(port/tunnel ID
10321 Ext-ID 192.168.0.5) Proto 0
Mar 11 14:05:30.179395 RSVP PSB E-D, allocating psb resources for label 4294967295
Mar 11 14:05:30.180353 LSP "E-D" is Up
Session: 192.168.0.4(port/tunnel ID 10321 Ext-ID 192.168.0.5) Proto 0
Sender: 192.168.0.5(port/lsp ID 2)

```



## CHAPTER 11

# Configuring RSVP Link Protection and Node Protection to Protect from Traffic Failures

- [Link Protection on page 495](#)
- [Multiple Bypass LSPs for Link Protection on page 496](#)
- [Node Protection on page 497](#)
- [Fast Reroute, Node Protection, and Link Protection on page 498](#)
- [Configuring Link Protection on Interfaces Used by LSPs on page 502](#)
- [Configuring Node Protection or Link Protection for LSPs on page 509](#)
- [Configuring Inter-AS Node and Link Protection on page 510](#)

## Link Protection


---

Link protection helps to ensure that traffic going over a specific interface to a neighboring router or switch can continue to reach this router (switch) if that interface fails. When link protection is configured for an interface and an LSP that traverses this interface, a bypass LSP is created that will handle this traffic if the interface fails. The bypass LSP uses a different interface and path to reach the same destination. The path used can be configured explicitly, or you can rely on CSPF. The RSVP metric for the bypass LSP is set in the range of 20,000 through 29,999 (this value is not user configurable).

If a link-protected interface fails, traffic is quickly switched to the bypass LSP. Note that a bypass LSP cannot share the same egress interface with the LSPs it monitors.

In [Figure 42 on page 496](#), link protection is enabled on Interface B between Router 1 and Router 2. It is also enabled on LSP A, an LSP that traverses the link between Router 1 and Router 2. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the bypass LSP generated by link protection.

**Figure 42: Link Protection Creating a Bypass LSP for the Protected Interface**



Although LSPs traversing an interface can be configured to take advantage of link protection, it is important to note that it is specifically the interface that benefits from link protection. If link protection is enabled on an interface but not on a particular LSP traversing that interface, then if the interface fails, that LSP will also fail.



**NOTE:** Link protection does not work on unnumbered interfaces.

To protect traffic over the entire route taken by an LSP, you should configure fast reroute. For more information, see [“Configuring Fast Reroute” on page 226](#).

**Related Documentation**

- [Node Protection on page 497](#)
- [Multiple Bypass LSPs for Link Protection on page 496](#)
- [Fast Reroute, Node Protection, and Link Protection on page 498](#)
- [Configuring Node Protection or Link Protection for LSPs on page 509](#)

## Multiple Bypass LSPs for Link Protection

By default, link protection relies on a single bypass LSP to provide path protection for an interface. However, you can also specify multiple bypass LSPs to provide link protection for an interface. You can individually configure each of these bypass LSPs or create a single configuration for all of the bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints.

The following algorithm describes how and when an additional bypass LSP is activated for an LSP:

1. If any currently active bypass can satisfy the requirements of the LSP (bandwidth, link protection, or node-link protection), the traffic is directed to that bypass.
2. If no active bypass LSP is available, scan through the manual bypass LSPs in first-in, first-out (FIFO) order, skipping those that are already active (each manual bypass can only be activated once). The first inactive manual bypass that can satisfy the requirements is activated and traffic is directed to that bypass.
3. If no manual bypass LSPs are available and if the **max-bypasses** statement activates multiple bypass LSPs for link protection, determine whether an automatically configured bypass LSP can satisfy the requirements. If an automatically configured

bypass LSP is available and if the total number of active automatically configured bypass LSPs does not exceed the maximum bypass LSP limit (configured with the **max-bypasses** statement), activate another bypass LSP.

For information about how to configure multiple bypass LSPs for link protection, see [“Configuring Bypass LSPs” on page 503](#).

## Node Protection


Node protection extends the capabilities of link protection. Link protection helps to ensure that traffic going over a specific interface to a neighboring router can continue to reach this router if that interface fails. Node protection ensures that traffic from an LSP traversing a neighboring router can continue to reach its destination even if the neighboring router fails.

When you enable node protection for an LSP, you must also enable link protection. Once enabled, node protection and link protection establish the following types of bypass LSPs:

- Next-hop bypass LSP—Provides an alternate route for an LSP to reach a neighboring router. This type of bypass LSP is established when you enable either node protection or link protection.
- Next-next-hop bypass LSP—Provides an alternate route for an LSP to get around a neighboring router en route to the destination router. This type of bypass LSP is established exclusively when node protection is configured. If a next-next-hop bypass LSP cannot be created, an attempt is made to signal a next-hop bypass LSP.

In [Figure 43 on page 497](#), node protection is enabled on Interface B on Router 1. Node protection is also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.

**Figure 43: Node Protection Creating a Next-Next-Hop Bypass LSP**



The time needed by node protection to switch traffic to a next-next-hop bypass LSP can be significantly longer than the time needed by link protection to switch traffic to a next-hop bypass LSP. Link protection relies on a hardware mechanism to detect a link failure, allowing it to quickly switch traffic to a next-hop bypass LSP.

Node failures are often due to software problems on the node router. Node protection relies on the receipt of hello messages from a neighboring router to determine whether it is still functioning. The time it takes node protection to divert traffic partly depends on how often the node router sends hello messages and how long it takes the node-protected

router to react to having not received a hello message. However, once the failure is detected, traffic can be quickly diverted to the next-next-hop bypass LSP.



**NOTE:**

Node protection provides traffic protection in the event of an error or interruption of the physical link between two routers. It does not provide protection in the event of control plane errors. The following provides an example of a control plane error:

- A transit router changes the label of a packet due to a control plane error.
- When the ingress router receives the packet, it considers the label change to be a catastrophic event and deletes both the primary LSP and the associated bypass LSP.

**Related Documentation**

- [Link Protection on page 495](#)
- [Fast Reroute, Node Protection, and Link Protection on page 498](#)
- [Configuring Node Protection or Link Protection for LSPs on page 509](#)

---

## Fast Reroute, Node Protection, and Link Protection

This document discusses the following sections:

- [LSP Protection Overview on page 498](#)
- [LSP Protection Types Comparison on page 499](#)
- [One-to-One Backup Implementation on page 499](#)
- [Facility Backup Implementation on page 500](#)

### LSP Protection Overview

RSVP-TE extensions establish backup label-switched path (LSP) tunnels for local repair of LSP tunnels. These mechanisms enable immediate re-direction of traffic onto backup LSP tunnels, in the event of a failure.

RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*, describes two different types of traffic protection for RSVP-signaled LSPs:

- One-to-one backup—In this method, detour LSPs for each protected LSP is created at each potential point of local repair.
- Facility backup—In this method, a bypass tunnel is created to protect a set of LSPs that have similar backup constraints at a potential failure point, by taking advantage of the MPLS label stacking.

The one-to-one backup and the facility backup methods protect links and nodes during network failure, and can co-exist in a mixed network.



## LSP Protection Types Comparison

In the Junos OS, the one-to-one backup of traffic protection is provided by fast reroute. Each LSP requires a protecting LSP to be signaled at each hop except the egress router. This method of LSP protection cannot be shared.

In the facility backup method, the LSP traffic protection is provided on the node and link. Each LSP requires a protecting LSP to be signaled at each hop except the egress router. Unlike fast reroute, this protecting LSP can be shared by other LSPs.

Table 14 on page 499 summarizes the traffic protection types.

**Table 14: One-to-One Backup Compared with Facility Backup**

| Comparison                     | One-to-One Backup   | Facility Backup                                        |
|--------------------------------|---------------------|--------------------------------------------------------|
| Name of the protecting LSP     | Detour LSP          | Bypass LSP                                             |
| Sharing of the protecting LSP  | Cannot be shared    | Can be shared by multiple LSPs                         |
| Junos configuration statements | <b>fast-reroute</b> | <b>node-link-protection</b> and <b>link-protection</b> |

## One-to-One Backup Implementation

In the one-to-one backup method, the points of local repair maintain separate backup paths for each LSP passing through a facility. The backup path terminates by merging back with the primary path at a node called the merge point. In this approach, the merge point can be any node downstream from the protected facility.


In the one-to-one backup method, an LSP is established that intersects the original LSP downstream of the point of link or node failure. A separate backup LSP is established for each LSP that is backed up.

One-to-one backup is appropriate under the following circumstances:

- Protection of a small number of LSPs relative to the total number of LSPs.
- Path selection criteria, such as bandwidth, priority, and link coloring for detour paths is critical.
- Control of individual LSPs is important.

In Figure 44 on page 500, Routers R1 and R5 are the ingress and egress routers, respectively. A protected LSP is established between the two routers transiting Routers R2, R3, and R4. Router R2 provides user traffic protection by creating a partial backup LSP that merges with the protected LSP at Router R4. This partial one-to-one backup LSP is called a detour. Detours are always calculated to avoid the immediate downstream link and node, providing against both link and node failure.

Figure 44: One-to-One Backup



In the example, the protected LSP is **R1-R2-R3-R4-R5**, and the following detours are established:

- Router R1—**R1-R6-R7-R8-R3**
- Router R2—**R2-R7-R8-R4**
- Router R3—**R3-R8-R9-R5**
- Router R4—**R4-R9-R5**

To protect an LSP that traverses **N** nodes fully, there can be as many as **(N - 1)** detours. The point of local repair sends periodic refresh messages to maintain each backup path, as a result maintaining state information for backup paths protecting individual LSPs is a significant resource burden for the point of local repair. To minimize the number of LSPs in the network, it is desirable to merge a detour back to its protected LSP, when feasible. When a detour LSP intersects its protected LSP at an LSR with the same outgoing interface, it is merged.

## Facility Backup Implementation

In the facility backup approach, a point of local repair maintains a single backup path to protect a set of primary LSPs traversing the point of local repair, the facility, and the merge point. The facility backup is based on interface rather than on LSP. While fast reroute protects interfaces or nodes along the entire path of a LSP, the facility backup protection can be applied on interfaces as needed. As a result, fewer states need to be maintained and refreshed which results in a scalable solution. The facility backup method is also called many-to-one backup.

The facility backup method takes advantage of the MPLS label stack. Instead of creating a separate LSP for every backed-up LSP, a single LSP is created that serves to back up a set of LSPs. Such an LSP tunnel is called a bypass tunnel. In this method, a router immediately upstream from a link failure uses an alternate interface to forward traffic

to its downstream neighbor, and the merge point should be the node immediately downstream to the facility. This is accomplished by preestablishing a bypass path that is shared by all protected LSPs traversing the failed link. A single bypass path can safeguard a set of protected LSPs. When an outage occurs, the router immediately upstream from the link outage switches protected traffic to the bypass link, then signals the link failure to the ingress router.


The bypass tunnel must intersect the path of the original LSP(s) somewhere downstream of the point of local repair. This constrains the set of LSPs being backed up through that bypass tunnel to those that pass through some common downstream nodes. All LSPs that pass through the point of local repair and through this common node, and that do not also use the facilities involved in the bypass tunnel are candidates for this set of LSPs.

The facility backup method is appropriate in the following situations:

- The number of LSPs to be protected is large.
- Satisfying path selection criteria (priority, bandwidth, and link coloring) for bypass paths is less critical.
- Control at the granularity of individual LSPs is not required.

In [Figure 45 on page 501](#), Routers R1 and R5 are the ingress and egress routers, respectively. Router R2 has established a bypass tunnel that protects against the failure of Router R2-R3 link and Router R3 node. A bypass tunnel is established between Routers R6 and R7. There are three different protected LSPs that are using the same bypass tunnel for protection.

**Figure 45: Facility Backup**



8042700

The facility backup method provides a scalability improvement, wherein the same bypass tunnel is also used to protect LSPs from any of Routers R1, R2, or R8 to any of Routers R4, R5, or R9.

**Related  
Documentation**

- [Fast Reroute Overview on page 46](#)
- [Configuring Fast Reroute on page 226](#)
- [Node Protection on page 497](#)
- [Configuring Node Protection or Link Protection for LSPs on page 509](#)

---

## Configuring Link Protection on Interfaces Used by LSPs

When you configure node protection or link protection on a router for LSPs as described in “[Configuring Node Protection or Link Protection for LSPs](#)” on [page 509](#), you also must configure the **link-protection** statement on the RSVP interfaces used by the LSPs.

To configure link protection on the interfaces used by the LSPs, include the **link-protection** statement:

```
link-protection {
 disable;
 admin-group
 exclude group-names;
 include-all group-names;
 include-any group-names;
}
bandwidth bps;
bypass bypass-name {
 bandwidth bps;
 description text;
 hop-limit number;
 no-cspf;
 path address <strict | loose>;
 priority setup-priority reservation-priority;
 to address;
}
class-of-service cos-value;
hop-limit number;
max-bypasses number;
no-cspf;
no-node-protection;
optimize-timer seconds;
path address <strict | loose>;
priority setup-priority reservation-priority;
subscription percent {
 ct0 percent;
 ct1 percent;
 ct2 percent;
 ct3 percent;
}
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name*]

All the statements under **link-protection** are optional.

The following sections describe how to configure link protection:

- [Configuring Bypass LSPs on page 503](#)
- [Configuring Administrative Groups for Bypass LSPs on page 504](#)
- [Configuring the Bandwidth for Bypass LSPs on page 504](#)
- [Configuring Class of Service for Bypass LSPs on page 505](#)
- [Configuring the Hop Limit for Bypass LSPs on page 505](#)
- [Configuring the Maximum Number of Bypass LSPs on page 506](#)
- [Disabling CSPF for Bypass LSPs on page 506](#)
- [Disabling Node Protection for Bypass LSPs on page 507](#)
- [Configuring the Optimization Interval for Bypass LSPs on page 507](#)
- [Configuring an Explicit Path for Bypass LSPs on page 508](#)
- [Configuring the Amount of Bandwidth Subscribed for Bypass LSPs on page 508](#)
- [Configuring Priority and Preemption for Bypass LSPs on page 509](#)

## Configuring Bypass LSPs

You can configure specific bandwidth and path constraints for a bypass LSP. You can also individually configure each bypass LSP generated when you enable multiple bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints (if any).

If you specify the **bandwidth**, **hop-limit**, and **path** statements for the bypass LSP, these values take precedence over the values configured at the [edit protocols rsvp **interface** *interface-name* **link-protection**] hierarchy level. The other attributes (**subscription**, **no-node-protection**, and **optimize-timer**) are inherited from the general constraints.

To configure a bypass LSP, specify a name for the bypass LSP using the **bypass** statement. The name can be up to 64 characters in length.

```
bypass bypass-name {
 bandwidth bps;
 description text;
 class-of-service cos-value;
 hop-limit number;
 no-cspf;
 path address <strict | loose>;
 priority setup-priority reservation-priority;
 to address;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* **link-protection**]

- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

### Configuring the Next-Hop or Next-Next-Hop Node Address for Bypass LSPs

---

If you configure a bypass LSP, you must also configure the **to** statement. The **to** statement specifies the address for the interface of the immediate next-hop node (for link protection) or the next-next-hop node (for node-link protection). The address specified determines whether this is a link protection bypass or a node-link protection bypass. On multiaccess networks (for example, a LAN), this address is also used to specify which next-hop node is being protected.

### Configuring Administrative Groups for Bypass LSPs

Administrative groups, also known as link coloring or resource class, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. You can use administrative groups to implement a variety of policy-based LSP setups. You can configure administrative groups for bypass LSPs. For more information about configuring administrative groups, see “[Configuring Administrative Groups for LSPs](#)” on page 240.

To configure administrative groups for bypass LSPs, include the **admin-group** statement:

```
admin-group {
 exclude group-names;
 include-all group-names;
 include-any group-names;
}
```

To configure an administrative group for all of the bypass LSPs, include the **admin-group** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

To configure an administrative groups for a specific bypass LSP, include the **admin-group** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

### Configuring the Bandwidth for Bypass LSPs

You can specify the amount of bandwidth allocated for automatically generated bypass LSPs or you can individually specify the amount of bandwidth allocated for each LSP.

If you have enabled multiple bypass LSPs, this statement is required.

To specify the bandwidth allocation, include the **bandwidth** statement:

**bandwidth** *bps*;

For automatically generated bypass LSPs, include the **bandwidth** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

For individually configured bypass LSPs, include the **bandwidth** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

## Configuring Class of Service for Bypass LSPs

You can specify the class-of-service value for bypass LSPs by including the **class-of-service** statement:

**class-of-service** *cos-value*;

To apply a class-of-service value to all the automatically generated bypass LSPs, include the **class-of-service** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

To configure a class-of-service value for a specific bypass LSPs, include the **class-of-service** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

## Configuring the Hop Limit for Bypass LSPs

You can specify the maximum number of hops a bypass can traverse. By default, each bypass can traverse a maximum of 255 hops (the ingress and egress routers count as one hop each, so the minimum hop limit is two).

To configure the hop limit for bypass LSPs, include the **hop-limit** statement:

**hop-limit** *number*;

For automatically generated bypass LSPs, include the **hop-limit** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]

- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

For individually configured bypass LSPs, include the **hop-limit** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

## Configuring the Maximum Number of Bypass LSPs

You can specify the maximum number of dynamic bypass LSPs permitted for protecting an interface using the **max-bypasses** statement at the [edit protocols rsvp **interface** *interface-name* link-protection] hierarchy level. When this statement is configured, multiple bypasses for link protection are enabled. Call admission control (CAC) is also enabled.

By default, this option is disabled and only one bypass is enabled for each interface. You can configure a value of between 0 through 99 for the **max-bypasses** statement. Configuring a value of 0 prevents the creation of any dynamic bypass LSPs for the interface. If you configure a value of 0 for the **max-bypasses** statement, you need to configure one or more static bypass LSPs to enable link protection on the interface.

If you configure the **max-bypasses** statement, you must also configure the **bandwidth** statement (discussed in “[Configuring the Bandwidth for Bypass LSPs](#)” on page 504).

To configure the maximum number of bypass LSPs for a protected interface, include the **max-bypasses** statement:

```
max-bypasses number;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

## Disabling CSPF for Bypass LSPs

Under certain circumstances, you might need to disable CSPF computation for bypass LSPs and use the configured Explicit Route Object (ERO) if available. For example, a bypass LSP might need to traverse multiple OSPF areas or IS-IS levels, preventing the CSPF computation from working. To ensure that link and node protection function properly in this case, you have to disable CSPF computation for the bypass LSP.

You can disable CSPF computation for all bypass LSPs or for specific bypass LSPs.

To disable CSPF computation for bypass LSPs, include the **no-cspf** statement:

```
no-cspf;
```

For a list of hierarchy levels where you can include this statement, see the statement summary for this statement.



## Disabling Node Protection for Bypass LSPs

You can disable node protection on the RSVP interface. Link protection remains active. When this option is configured, the router can only initiate a next-hop bypass, not a next-next-hop bypass.

To disable node protection for bypass LSPs, include the **no-node-protection** statement:

```
no-node-protection;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

## Configuring the Optimization Interval for Bypass LSPs

You can configure an optimization interval for bypass LSPs using the **optimize-timer** statement. At the end of this interval, an optimization process is initiated that attempts to either minimize the number of bypasses currently in use, minimize the total amount of bandwidth reserved for all of the bypasses, or both. You can configure an optimization interval from 1 through 65,535 seconds. A default value of 0 disables bypass LSP optimization.

When you configure the **optimize-timer** statement, bypass LSPs are reoptimized automatically when you configure or change the configuration of any of the following:

- Administrative group for a bypass LSP—The configuration for an administrative group has been changed on a link along the path used by the bypass LSP. Configure an administrative group using the **admin-group** statement at the [edit protocols rsvp **interface** *interface-name* link-protection] hierarchy level.
- Fate sharing group—The configuration for a fate sharing group has been changed. Configure a fate sharing group using the **group** statement at the [edit routing-options **fate-sharing**] hierarchy level.
- IS-IS overload—The configuration for IS-IS overload has been changed on a router along the path used by the bypass LSP. Configure IS-IS overload using the **overload** statement at the [edit protocols isis] hierarchy level.
- IGP metric—The IGP metric has been changed on a link along the path used by the bypass LSP.

To configure the optimization interval for bypass LSPs, include the **optimize-timer** statement:

```
optimize-timer seconds;
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]

- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

## Configuring an Explicit Path for Bypass LSPs

By default, when you establish a bypass LSP to an adjacent neighbor, CSPF is used to discover the least-cost path. The **path** statement allows you to configure an explicit path (a sequence of strict or loose routes), giving you control over where and how the bypass LSP is established. To configure an explicit path, include the **path** statement:

```
path address <strict | loose>;
```

For automatically generated bypass LSPs, include the **path** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

For individually configured bypass LSPs, include the **path** statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]
- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection **bypass** *bypass-name*]

## Configuring the Amount of Bandwidth Subscribed for Bypass LSPs

You can configure the amount of bandwidth subscribed to bypass LSPs. You can configure the bandwidth subscription for the whole bypass LSP or for each class type that might traverse the bypass LSP. You can configure any value between 1 percent and 65,535 percent. By configuring a value less than 100 percent, you are undersubscribing the bypass LSPs. By configuring a value greater than 100 percent, you are oversubscribing the bypass LSPs.

The ability to oversubscribe the bandwidth for the bypass LSPs makes it possible to more efficiently use network resources. You can configure the bandwidth for the bypass LSPs based on the average network load as opposed to the peak load.

To configure the amount of bandwidth subscribed for bypass LSPs, include the **subscription** statement:

```
subscription percentage {
 ct0 percentage;
 ct1 percentage;
 ct2 percentage;
 ct3 percentage;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols rsvp **interface** *interface-name* link-protection]

- [edit logical-systems *logical-system-name* protocols rsvp **interface** *interface-name* link-protection]

## Configuring Priority and Preemption for Bypass LSPs

When there is insufficient bandwidth to establish a more important LSP, you might want to tear down a less important existing LSP to release the bandwidth. You do this by preempting the existing LSP.

For more detailed information on configuring setup priority and reservation priority for LSPs, see “[Configuring Priority and Preemption for LSPs](#)” on page 250.

To configure the bypass LSP's priority and preemption properties, include the **priority** statement:

```
priority setup-priority reservation-priority;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

---

## Configuring Node Protection or Link Protection for LSPs

When you configure node protection or link protection on a router or switch, bypass LSPs are created to the next-hop or next-next-hop routers (switches) for the LSPs traversing the router (switch). You must configure node protection or link protection for each LSP that you want protected. To extend protection along the entire path used by an LSP, you must configure protection on each router that the LSP traverses.

You can configure node protection or link protection for both static and dynamic LSPs.

To configure node protection on a router for a specified LSP, include the **node-link-protection** statement:

```
node-link-protection;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]

To configure link protection on a router for a specified LSP, include the **link-protection** statement:

```
link-protection;
```

You can include this statement at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name*]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name*]



**NOTE:** To complete the configuration of node or link protection, you must also configure link protection on all unidirectional RSVP interfaces that the LSPs traverse, as described in [“Configuring Link Protection on Interfaces Used by LSPs” on page 502](#).

**Related  
Documentation**

- [Configuring Link Protection on Interfaces Used by LSPs on page 502](#)
- [Link Protection on page 495](#)
- [Node Protection on page 497](#)
- [Multiple Bypass LSPs for Link Protection on page 496](#)
- [Fast Reroute, Node Protection, and Link Protection on page 498](#)

---

## Configuring Inter-AS Node and Link Protection

To interoperate with other vendors' equipment, the Junos OS supports the record route object (RRO) node ID subobject for use in inter-AS link and node protection configurations. The RRO node ID subobject is defined in RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*. This functionality is enabled by default in Junos OS Release 9.4 and later.

If you have Juniper Networks routers running Junos OS Release 9.4 and later releases in the same MPLS-TE network as routers running Junos OS Release 8.4 and earlier releases, you might need to disable the RRO node ID subobject by configuring the **no-node-id-subobject** statement:

**no-node-id-subobject;**

You can include this statement at the following hierarchy levels:

- **[edit protocols rsvp]**
- **[edit logical-systems *logical-system-name* protocols rsvp]**

## CHAPTER 12

# Configuring RSVP Graceful Restart for High Availability

- [RSVP Graceful Restart on page 511](#)
- [RSVP Graceful Restart Standard on page 511](#)
- [RSVP Graceful Restart Terminology on page 512](#)
- [RSVP Graceful Restart Operation on page 512](#)
- [Processing the Restart Cap Object on page 513](#)
- [Configuring RSVP Graceful Restart on page 514](#)

## RSVP Graceful Restart

---

RSVP graceful restart allows a router undergoing a restart to inform its adjacent neighbors of its condition. The restarting router requests a grace period from the neighbor or peer, which can then cooperate with the restarting router. The restarting router can still forward MPLS traffic during the restart period; convergence in the network is not disrupted. The restart is not visible to the rest of the network, and the restarting router is not removed from the network topology. RSVP graceful restart can be enabled on both transit routers and ingress routers. It is available for both point-to-point LSPs and point-to-multipoint LSPs.

RSVP graceful restart is described in the following sections:

- [RSVP Graceful Restart Standard on page 511](#)
- [RSVP Graceful Restart Terminology on page 512](#)
- [RSVP Graceful Restart Operation on page 512](#)
- [Processing the Restart Cap Object on page 513](#)

## RSVP Graceful Restart Standard

---

RSVP graceful restart is described in RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions* (only Section 9, “Fault Handling”).

## RSVP Graceful Restart Terminology

---

### R

|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Recovery time<br/>(in milliseconds)</b> | <p>Applies only when the control channel is up (the hello exchange is complete) before the restart time. Applies only to nodal faults.</p> <p>When a graceful restart is in progress, the time left to complete a recovery is advertised. At other times, this value is zero. The maximum advertised recovery time is 2 minutes (120,000 milliseconds).</p> <p>During the recovery time, a restarting node attempts to recover its lost states with assistance from its neighbors. The neighbor of the restarting node must send the path messages with the recovery labels to the restarting node within a period of one-half the recovery time. The restarting node considers its graceful restart complete after its advertised recovery time.</p> |
| <b>Restart time<br/>(in milliseconds)</b>  | <p>The default value is 60,000 milliseconds (1 minute). The restart time is advertised in the hello message. The time indicates how long a neighbor should wait to receive a hello message from a restarting router before declaring that router dead and purging states.</p> <p>The Junos OS can override a neighbor's advertised restart time if the time is greater than one-third the local restart time. For example, given the default restart time of 60 seconds, a router would wait 20 seconds or less to receive a hello message from a restarting neighbor. If the restart time is zero, the restarting neighbor can immediately be declared dead.</p>                                                                                     |

## RSVP Graceful Restart Operation

---

For RSVP graceful restart to function, the feature must be enabled on the global routing instance. RSVP graceful restart can be disabled at the protocol level (for RSVP alone) or at the global level for all protocols.

RSVP graceful restart requires the following of a restarting router and the router's neighbors:

- For the restarting router, RSVP graceful restart attempts to maintain the routes installed by RSVP and the allocated labels, so that traffic continues to be forwarded without disruption. RSVP graceful restart is done quickly enough to reduce or eliminate the impact on neighboring nodes.
- The neighboring routers must have RSVP graceful restart helper mode enabled, thus allowing them to assist a router attempting to restart RSVP.

An object called Restart Cap that is sent in RSVP hello messages advertises a node's restart capability. The neighboring node sends a Recover Label object to the restarting node to recover its forwarding state. This object is essentially the old label that the restarting node advertised before the node went down.

The following lists the RSVP graceful restart behaviors, which vary depending on the configuration and on which features are enabled:

- If you disable helper mode, the Junos OS does not attempt to help a neighbor restart RSVP. Any information that arrives with a Restart Cap object from a neighbor is ignored.
- When you enable graceful restart under the routing instance configuration, the router can restart gracefully with the help of its neighbors. RSVP advertises a Restart Cap object (RSVP RESTART) in hello messages in which restart and recovery times are specified (neither value is 0).
- If you explicitly disable RSVP graceful restart under the **[protocols rsvp]** hierarchy level, the Restart Cap object is advertised with restart and recovery times specified as 0. The restart of neighboring routers is supported (unless helper mode is disabled), but the router itself does not preserve the RSVP forwarding state and cannot recover its control state.
- If after a restart RSVP realizes that no forwarding state has been preserved, the Restart Cap object is advertised with restart and recovery times specified as 0.
- If graceful restart and helper mode are disabled, RSVP graceful restart is completely disabled. The router neither recognizes nor advertises the RSVP graceful restart objects.

You cannot explicitly configure values for the restart and recovery times.

Unlike other protocols, there is no way for RSVP to determine that it has completed a restart procedure, other than a fixed timeout. All RSVP graceful restart procedures are timer-based. A **show rsvp version** command might indicate that the restart is still in progress even if all RSVP sessions are back up and the routes are restored.

## Processing the Restart Cap Object

---

The following assumptions are made about a neighbor based on the Restart Cap object (assuming that a control channel failure can be distinguished unambiguously from a node restart):

- A neighbor that does not advertise the Restart Cap object in its hello messages cannot assist a router with state or label recovery, nor can it perform an RSVP graceful restart.
- After a restart, a neighbor advertising a Restart Cap object with a restart time equal to any value and a recovery time equal to 0 has not preserved its forwarding state. When a recovery time equals 0, the neighbor is considered dead and any states related to this neighbor are purged, regardless of the value of the restart time.
- After a restart, a neighbor advertising its recovery time with a value other than 0 can keep or has kept the forwarding state. If the local router is helping its neighbor with restart or recovery procedures, it sends a Recover Label object to this neighbor.

## Configuring RSVP Graceful Restart

---

The following RSVP graceful restart configurations are possible:

- Graceful restart and helper mode are both enabled (the default).
- Graceful restart is enabled but helper mode is disabled. A router configured in this way can restart gracefully, but cannot help a neighbor with its restart and recovery procedures.
- Graceful restart is disabled but helper mode is enabled. A router configured in this way cannot restart gracefully, but can help a restarting neighbor.
- Graceful restart and helper mode both are disabled. This configuration completely disables RSVP graceful restart (including restart and recovery procedures and helper mode). The router behaves like a router that does not support RSVP graceful restart.



**NOTE:** In order to turn on RSVP graceful restart, you must set the global graceful restart timer to at least 180 seconds.

The following sections describe how to configure RSVP graceful restart:

- [Enabling Graceful Restart for All Routing Protocols on page 514](#)
- [Disabling Graceful Restart for RSVP on page 514](#)
- [Disabling RSVP Helper Mode on page 515](#)
- [Configuring the Maximum Helper Recovery Time on page 515](#)
- [Configuring the Maximum Helper Restart Time on page 515](#)

### Enabling Graceful Restart for All Routing Protocols

To enable graceful restart for RSVP, you need to enable graceful restart for all the protocols that support graceful restart on the router. For more information about graceful restart, see the *Junos OS Routing Protocols Library for Routing Devices*.

To enable graceful restart on the router, include the **graceful-restart** statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options]**
- **[edit logical-systems *logical-system-name* routing-options]**

### Disabling Graceful Restart for RSVP

By default, RSVP graceful restart and RSVP helper mode are enabled when you enable graceful restart. However, you can disable one or both of these capabilities.



To disable RSVP graceful restart and recovery, include the **disable** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level:

```
disable;
```

## Disabling RSVP Helper Mode

To disable RSVP helper mode, include the **helper-disable** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level:

```
helper-disable;
```

## Configuring the Maximum Helper Recovery Time

To configure the amount of time the router retains the state of its RSVP neighbors while they undergo a graceful restart, include the **maximum-helper-recovery-time** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to recover.

```
maximum-helper-recovery-time seconds;
```

## Configuring the Maximum Helper Restart Time

To configure the delay between when the router discovers that a neighboring router has gone down and when it declares the neighbor down, include the **maximum-helper-restart-time** statement at the **[edit protocols rsvp graceful-restart]** hierarchy level. This value is applied to all neighboring routers, so it should be based on the time required by the slowest RSVP neighbor to restart.

```
maximum-helper-restart-time seconds;
```



## PART 4

# Configuring LDP

- [LDP Overview on page 519](#)
- [Configuring LDP on page 527](#)



## CHAPTER 13

# LDP Overview

- [LDP Introduction on page 519](#)
- [Supported LDP Standards on page 520](#)
- [Junos OS LDP Protocol Implementation on page 520](#)
- [LDP Operation on page 521](#)
- [LDP Message Types on page 521](#)
- [Discovery Messages on page 521](#)
- [Session Messages on page 522](#)
- [Advertisement Messages on page 522](#)
- [Notification Messages on page 522](#)
- [Tunneling LDP LSPs in RSVP LSPs on page 523](#)
- [Tunneling LDP LSPs in RSVP LSPs Overview on page 523](#)
- [Label Operations on page 523](#)
- [LDP Session Protection on page 525](#)

## LDP Introduction

---

The Label Distribution Protocol (LDP) is a protocol for distributing 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.

These LSPs might have an endpoint at a directly attached neighbor (comparable to IP hop-by-hop forwarding), or at a network egress node, enabling switching through all intermediary nodes. LSPs established by LDP can also traverse traffic-engineered LSPs created by RSVP.

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers. This process forms a tree of LSPs that converge on the egress router.

## Supported LDP Standards

---

Junos OS substantially supports the following RFCs and Internet drafts, which define standards for LDP.

- RFC 3212, *Constraint-Based LSP Setup using LDP*
- RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol*
- Internet draft draft-napierala-mpls-targeted-mldp-01.txt, *Using LDP Multipoint Extensions on Targeted LDP Sessions*

The following RFCs do not define standards, but provide information about LDP. The IETF classifies them as “Informational.”

- RFC 3215, *LDP State Machine*
- RFC 5036, *LDP Specification*

For the following features described in the indicated sections of the RFC, Junos OS supports one of the possible modes but not the others:

- Label distribution control (section 2.6.1): Ordered mode is supported, but not Independent mode.
- Label retention (section 2.6.2): Liberal mode is supported, but not Conservative mode.
- Label advertisement (section 2.6.3): Both Downstream Unsolicited mode and Downstream on Demand mode are supported.
- RFC 5443, *LDP IGP Synchronization*
- RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*

Junos OS support limited to point-to-multipoint extensions for LDP.

### Related Documentation

- [Supported GMPLS Standards on page 673](#)
- [Supported MPLS Standards on page 20](#)
- [Supported RSVP Standards on page 460](#)
- *Accessing Standards Documents on the Internet*

## Junos OS LDP Protocol Implementation

---

The Junos OS implementation of LDP supports LDP version 1. The Junos OS supports a simple mechanism for tunneling between routers in an interior gateway protocol (IGP), to eliminate the required distribution of external routes within the core. The Junos OS allows an MPLS tunnel next hop to all egress routers in the network, with only an IGP running in the core to distribute routes to egress routers. Edge routers run BGP but do not distribute external routes to the core. Instead, the recursive route lookup at the edge

resolves to an LSP switched to the egress router. No external routes are necessary on the transit LDP routers.

## LDP Operation

You must configure LDP for each interface on which you want LDP to run. LDP creates LSP trees rooted at each egress router for the router ID address that is the subsequent BGP next hop. The ingress point is at every router running LDP. This process provides an inet.3 route to every egress router. If BGP is running, it will attempt to resolve next hops by using the inet.3 table first, which binds most, if not all, of the BGP routes to MPLS tunnel next hops.

Two adjacent routers running LDP become neighbors. If the two routers are connected by more than one interface, they become neighbors on each interface. When LDP routers become neighbors, they establish an LDP session to exchange label information. If per-router labels are in use on both routers, only one LDP session is established between them, even if they are neighbors on multiple interfaces. For this reason, an LDP session is not related to a particular interface.

LDP operates in conjunction with a unicast routing protocol. LDP installs LSPs only when both LDP and the routing protocol are enabled. For this reason, you must enable both LDP and the routing protocol on the same set of interfaces. If this is not done, LSPs might not be established between each egress router and all ingress routers, which might result in loss of BGP-routed traffic.

You can apply policy filters to labels received from and distributed to other routers through LDP. Policy filters provide you with a mechanism to control the establishment of LSPs.

For LDP to run on an interface, MPLS must be enabled on a logical interface on that interface. For more information, see the *Logical Interfaces*.

**Related Documentation**

- [Logical Interfaces](#)

## LDP Message Types

LDP uses the message types described in the following sections to establish and remove mappings and to report errors. All LDP messages have a common structure that uses a type, length, and value (TLV) encoding scheme.

- [Discovery Messages on page 521](#)
- [Session Messages on page 522](#)
- [Advertisement Messages on page 522](#)
- [Notification Messages on page 522](#)

## Discovery Messages

Discovery messages announce and maintain the presence of a router in a network. Routers indicate their presence in a network by sending hello messages periodically. Hello

messages are transmitted as UDP packets to the LDP port at the group multicast address for all routers on the subnet.

LDP uses the following discovery procedures:

- Basic discovery—A router periodically sends LDP link hello messages through an interface. LDP link hello messages are sent as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.
- Extended discovery—LDP sessions between routers not directly connected are supported by LDP extended discovery. A router periodically sends LDP targeted hello messages to a specific address. Targeted hello messages are sent as UDP packets addressed to the LDP discovery port at the specific address. The targeted router decides whether to respond to or ignore the targeted hello message. A targeted router that chooses to respond does so by periodically sending targeted hello messages to the initiating router.

---

## Session Messages

Session messages establish, maintain, and terminate sessions between LDP peers. When a router establishes a session with another router learned through the hello message, it uses the LDP initialization procedure over TCP transport. When the initialization procedure is completed successfully, the two routers are LDP peers and can exchange advertisement messages.

---

## Advertisement Messages

Advertisement messages create, change, and delete label mappings for forwarding equivalence classes (FECs). Requesting a label or advertising a label mapping to a peer is a decision made by the local router. In general, the router requests a label mapping from a neighboring router when it needs one and advertises a label mapping to a neighboring router when it wants the neighbor to use a label.

---

## Notification Messages

Notification messages provide advisory information and signal error information. LDP sends notification messages to report errors and other events of interest. There are two kinds of LDP notification messages:

- Error notifications, which signal fatal errors. If a router receives an error notification from a peer for an LDP session, it terminates the LDP session by closing the TCP transport connection for the session and discarding all label mappings learned through the session.
- Advisory notifications, which pass information to a router about the LDP session or the status of some previous message received from the peer.



## Tunneling LDP LSPs in RSVP LSPs

You can tunnel LDP LSPs over RSVP LSPs. The following sections describe how tunneling of LDP LSPs in RSVP LSPs works:

- [Tunneling LDP LSPs in RSVP LSPs Overview on page 523](#)
- [Label Operations on page 523](#)

## Tunneling LDP LSPs in RSVP LSPs Overview

If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.

When you configure the router to run LDP across RSVP-established LSPs, LDP automatically establishes sessions with the router at the other end of the LSP. LDP control packets are routed hop-by-hop, rather than carried through the LSP. This routing allows you to use simplex (one-way) traffic-engineered LSPs. Traffic in the opposite direction flows through LDP-established LSPs that follow unicast routing rather than through traffic-engineered tunnels.

If you configure LDP over RSVP LSPs, you can still configure multiple OSPF areas and IS-IS levels in the traffic engineered core and in the surrounding LDP cloud.



**NOTE:** Beginning with Junos OS Release 15.1, multi-instance support is extended to LDP over RSVP tunneling for a virtual router routing instance. This allows splitting of a single routing and MPLS domain into multiple domains so that each domain can be scaled independently. BGP labeled unicast can be used to stitch these domains for service FECs. Each domain uses intra-domain LDP over RSVP LSP for MPLS forwarding.

### Related Documentation

- [Label Operations on page 523](#)
- [Configuring a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs Over a Single RSVP LSP on page 724](#)

## Label Operations

[Figure 46 on page 524](#) depicts an LDP LSP being tunneled through an RSVP LSP. (For definitions of label operations, see [“MPLS Label Overview” on page 24](#).) The shaded inner oval represents the RSVP domain, whereas the outer oval depicts the LDP domain. RSVP establishes an LSP through routers B, C, D, and E, with the sequence of labels L3, L4. LDP establishes an LSP through Routers A, B, E, F, and G, with the sequence of labels L1, L2, L5. LDP views the RSVP LSP between Routers B and E as a single hop.

When the packet arrives at Router A, it enters the LSP established by LDP, and a label (L1) is pushed onto the packet. When the packet arrives at Router B, the label (L1) is swapped with another label (L2). Because the packet is entering the traffic-engineered LSP established by RSVP, a second label (L3) is pushed onto the packet.

This outer label (L3) is swapped with a new label (L4) at the intermediate router (C) within the RSVP LSP tunnel, and when the penultimate router (D) is reached, the top label is popped. Router E swaps the label (L2) with a new label (L5), and the penultimate router for the LDP-established LSP (F) pops the last label.

Figure 46: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs





Figure 47 on page 524 depicts a double push label operation (L1L2). A double push label operation is used when the ingress router (A) for both the LDP LSP and the RSVP LSP tunneled through it is the same device. Note that Router D is the penultimate hop for the LDP-established LSP, so L2 is popped from the packet by Router D.

Figure 47: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs



---

## LDP Session Protection

---

LDP session protection is based on the LDP targeted hello functionality defined in RFC 5036, *LDP Specification*, and is supported by the Junos OS as well as the LDP implementations of most other vendors. It involves sending unicast User Datagram Protocol (UDP) hello packets to a remote neighbor address and receiving similar packets from the neighbor router.

If you configure LDP session protection on a router, the LDP sessions are maintained as follows:

1. An LDP session is established between a router and a remote neighboring router.
2. If all of the direct links between the routers go down, the LDP session remains up so long as there is IP connectivity between the routers based on another connection over the network.
3. When the direct link between the routers is reestablished, the LDP session is not restarted. The routers simply exchange LDP hellos with each other over the direct link. They can then begin forwarding LDP-signaled MPLS packets using the original LDP session.

By default, LDP targeted hellos are set to the remote neighbor so long as the LDP session is up, even if there are no more link neighbors to that router. You can also specify the duration you would like to maintain the remote neighbor connection in the absence of link neighbors. When the last link neighbor for a session goes down, the Junos OS starts an LDP session protection timer. If this timer expires before any of the link neighbors come back up, the remote neighbor connection is taken down and the LDP session is terminated. If you configure a different value for the timer while it is currently running, the Junos OS updates the timer to the specified value without disrupting the current state of the LDP session.



## CHAPTER 14

# Configuring LDP

- [Minimum LDP Configuration on page 528](#)
- [Enabling and Disabling LDP on page 528](#)
- [Configuring the LDP Timer for Hello Messages on page 528](#)
- [Configuring the Delay Before LDP Neighbors Are Considered Down on page 529](#)
- [Enabling Strict Targeted Hello Messages for LDP on page 531](#)
- [Configuring the Interval for LDP Keepalive Messages on page 531](#)
- [Configuring the LDP Keepalive Timeout on page 531](#)
- [Configuring LDP Route Preferences on page 532](#)
- [LDP Graceful Restart on page 532](#)
- [Configuring LDP Graceful Restart on page 533](#)
- [Filtering Inbound LDP Label Bindings on page 535](#)
- [Filtering Outbound LDP Label Bindings on page 537](#)
- [Specifying the Transport Address Used by LDP on page 539](#)
- [Configuring the Prefixes Advertised into LDP from the Routing Table on page 540](#)
- [Configuring FEC Deaggregation on page 541](#)
- [Configuring Policers for LDP FECs on page 541](#)
- [Configuring LDP IPv4 FEC Filtering on page 542](#)
- [Configuring BFD for LDP LSPs on page 543](#)
- [Configuring ECMP-Aware BFD for LDP LSPs on page 546](#)
- [Configuring a Failure Action for the BFD Session on an LDP LSP on page 546](#)
- [Configuring the Holddown Interval for the BFD Session on page 547](#)
- [Configuring OAM Ingress Policies for LDP on page 547](#)
- [Configuring LDP Link Protection on page 548](#)
- [Example: Configuring LDP Link Protection on page 549](#)
- [Understanding Multicast-Only Fast Reroute on page 566](#)
- [Configuring Multicast-Only Fast Reroute on page 573](#)
- [Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain on page 576](#)

- [Example: Configuring LDP Downstream on Demand on page 592](#)
- [Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs on page 597](#)
- [Configuring Miscellaneous LDP Properties on page 628](#)
- [Configuring LDP LSP Traceroute on page 634](#)
- [Collecting LDP Statistics on page 635](#)
- [Tracing LDP Protocol Traffic on page 637](#)

---

## Minimum LDP Configuration

To enable LDP on a single interface, include the **ldp** statement and specify the interface using the **interface** statement. This is the minimum LDP configuration. All other LDP configuration statements are optional.

```
ldp {
 interface interface-name;
}
```

To enable LDP on all interfaces, specify **all** for *interface-name*.

For a list of hierarchy levels at which you can include these statements, see the statement summary sections.

---

## Enabling and Disabling LDP

LDP is routing-instance-aware. To enable LDP on a specific interface, include the following statements:

```
ldp {
 interface interface-name;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections.

To enable LDP on all interfaces, specify **all** for *interface-name*.

If you have configured interface properties on a group of interfaces and want to disable LDP on one of the interfaces, include the **interface** statement with the **disable** option:

```
interface interface-name {
 disable;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

---

## Configuring the LDP Timer for Hello Messages

LDP hello messages enable LDP nodes to discover one another and to detect the failure of a neighbor or the link to the neighbor. Hello messages are sent periodically on all interfaces where LDP is enabled.

There are two types of LDP hello messages:

- Link hello messages—Sent through the LDP interface as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.
- Targeted hello messages—Sent as UDP packets addressed to the LDP discovery port at a specific address. Targeted hello messages are used to support LDP sessions between routers that are not directly connected. A targeted router determines whether to respond or ignore a targeted hello message. A targeted router that chooses to respond does so by periodically sending targeted hello messages back to the initiating router.

By default, LDP sends hello messages every 5 seconds for link hello messages and every 15 seconds for targeted hello messages. You can configure the LDP timer to alter how often both types of hello messages are sent. However, you cannot configure a time for the LDP timer that is greater than the LDP hold time. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 529](#).

### Configuring the LDP Timer for Link Hello Messages

To modify how often LDP sends link hello messages, specify a new link hello message interval for the LDP timer using the **hello-interval** statement:

```
hello-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

### Configuring the LDP Timer for Targeted Hello Messages

To modify how often LDP sends targeted hello messages, specify a new targeted hello message interval for the LDP timer by configuring the **hello-interval** statement as an option for the **targeted-hello** statement:

```
targeted-hello {
 hello-interval seconds;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

---

### Configuring the Delay Before LDP Neighbors Are Considered Down

The hold time determines how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. The values sent by each neighbor do not have to match.

The hold time should normally be at least three times the hello interval. The default is 15 seconds for link hello messages and 45 seconds for targeted hello messages. However, it is possible to configure an LDP hold time that is close to the value for the hello interval.



**NOTE:** By configuring an LDP hold time close to the hello interval (less than three times the hello interval), LDP neighbor failures might be detected more quickly. However, this also increases the possibility that the router might declare an LDP neighbor down that is still functioning normally. For more information, see [“Configuring the LDP Timer for Hello Messages” on page 528](#).

The LDP hold time is also negotiated automatically between LDP peers. When two LDP peers advertise different LDP hold times to one another, the smaller value is used. If an LDP peer router advertises a shorter hold time than the value you have configured, the peer router's advertised hold time is used. This negotiation can affect the LDP keepalive interval as well.

If the local LDP hold time is not shortened during LDP peer negotiation, the user-configured keepalive interval is left unchanged. However, if the local hold time is reduced during peer negotiation, the keepalive interval is recalculated. If the LDP hold time has been reduced during peer negotiation, the keepalive interval is reduced to one-third of the new hold time value. For example, if the new hold-time value is 45 seconds, the keepalive interval is set to 15 seconds.

This automated keepalive interval calculation can cause different keepalive intervals to be configured on each peer router. This enables the routers to be flexible in how often they send keepalive messages, because the LDP peer negotiation ensures they are sent more frequently than the LDP hold time.

When you reconfigure the hold-time interval, changes do not take effect until after the session is reset. The hold time is negotiated when the LDP peering session is initiated and cannot be renegotiated as long as the session is up (required by RFC 5036, *LDP Specification*). To manually force the LDP session to reset, issue the **clear ldp session** command.

## Configuring the LDP Hold Time for Link Hello Messages

To modify how long an LDP node should wait for a link hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement:

```
hold-time seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Configuring the LDP Hold Time for Targeted Hello Messages

To modify how long an LDP node should wait for a targeted hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement as an option for the **targeted-hello** statement:

```
targeted-hello {
 hold-time seconds;
}
```



For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

## Enabling Strict Targeted Hello Messages for LDP

---

Use strict targeted hello messages to prevent LDP sessions from being established with remote neighbors that have not been specifically configured. If you configure the **strict-targeted-hellos** statement, an LDP peer does not respond to targeted hello messages coming from a source that is not one of its configured remote neighbors. Configured remote neighbors can include:

- Endpoints of RSVP tunnels for which LDP tunneling is configured
- Layer 2 circuit neighbors

If an unconfigured neighbor sends a hello message, the LDP peer ignores the message and logs an error (with the **error** trace flag) indicating the source. For example, if the LDP peer received a targeted hello from the Internet address 10.0.0.1 and no neighbor with this address is specifically configured, the following message is printed to the LDP log file:

LDP: Ignoring targeted hello from 10.0.0.1

To enable strict targeted hello messages, include the **strict-targeted-hellos** statement:

```
strict-targeted-hellos;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Configuring the Interval for LDP Keepalive Messages

---

The keepalive interval determines how often a message is sent over the session to ensure that the keepalive timeout is not exceeded. If no other LDP traffic is sent over the session in this much time, a keepalive message is sent. The default is 10 seconds. The minimum value is 1 second.

The value configured for the keepalive interval can be altered during LDP session negotiation if the value configured for the LDP hold time on the peer router is lower than the value configured locally. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 529](#).

To modify the keepalive interval, include the **keepalive-interval** statement:

```
keepalive-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Configuring the LDP Keepalive Timeout

---

After an LDP session is established, messages must be exchanged periodically to ensure that the session is still working. The keepalive timeout defines the amount of time that

the neighbor LDP node waits before deciding that the session has failed. This value is usually set to at least three times the keepalive interval. The default is 30 seconds.

To modify the keepalive interval, include the **keepalive-timeout** statement:

```
keepalive-timeout seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The value configured for the **keepalive-timeout** statement is displayed as the hold time when you issue the **show ldp session detail** command.

---

## Configuring LDP Route Preferences

When several protocols calculate routes to the same destination, route preferences are used to select which route is installed in the forwarding table. The route with the lowest preference value is selected. The preference value can be a number in the range 0 through 255. By default, LDP routes have a preference value of 9.

To modify the route preferences, include the **preference** statement:

```
preference preference;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

---

## LDP Graceful Restart

LDP graceful restart enables a router whose LDP control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. It also enables a router on which helper mode is enabled to assist a neighboring router that is attempting to restart LDP.

During session initialization, a router advertises its ability to perform LDP graceful restart or to take advantage of a neighbor performing LDP graceful restart by sending the graceful restart TLV. This TLV contains two fields relevant to LDP graceful restart: the reconnect time and the recovery time. The values of the reconnect and recovery times indicate the graceful restart capabilities supported by the router.

When a router discovers that a neighboring router is restarting, it waits until the end of the recovery time before attempting to reconnect. The recovery time is the length of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This time period is also typically the length of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

You can configure LDP graceful restart in both the master instance for the LDP protocol and for a specific routing instance. You can disable graceful restart at the global level for all protocols, at the protocol level for LDP only, and on a specific routing instance. LDP graceful restart is disabled by default, because at the global level, graceful restart

is disabled by default. However, helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default.

The following are some of the behaviors associated with LDP graceful restart:

- Outgoing labels are not maintained in restarts. New outgoing labels are allocated.
- When a router is restarting, no label-map messages are sent to neighbors that support graceful restart until the restarting router has stabilized (label-map messages are immediately sent to neighbors that do not support graceful restart). However, all other messages (keepalive, address-message, notification, and release) are sent as usual. Distributing these other messages prevents the router from distributing incomplete information.
- Helper mode and graceful restart are independent. You can disable graceful restart in the configuration, but still allow the router to cooperate with a neighbor attempting to restart gracefully.

## Configuring LDP Graceful Restart

When you alter the graceful restart configuration at either the **[edit routing-options graceful-restart]** or **[edit protocols ldp graceful-restart]** hierarchy levels, any running LDP session is automatically restarted to apply the graceful restart configuration. This behavior mirrors the behavior of BGP when you alter its graceful restart configuration.

By default, graceful restart helper mode is enabled, but graceful restart is disabled. Thus, the default behavior of a router is to assist neighboring routers attempting a graceful restart, but not to attempt a graceful restart itself.

To configure LDP graceful restart, see the following sections:

- [Enabling Graceful Restart on page 533](#)
- [Disabling LDP Graceful Restart or Helper Mode on page 534](#)
- [Configuring Reconnect Time on page 534](#)
- [Configuring Recovery Time and Maximum Recovery Time on page 535](#)

## Enabling Graceful Restart

To enable LDP graceful restart, you also need to enable graceful restart on the router.

To enable graceful restart, include the **graceful-restart** statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options]**
- **[edit logical-systems *logical-system-name* routing-options]**

The **graceful-restart** statement enables graceful restart for all protocols supporting this feature on the router. For more information about graceful restart, see the *Junos OS Routing Protocols Library for Routing Devices*.

By default, LDP graceful restart is enabled when you enable graceful restart at both the LDP protocol level and on all the routing instances. However, you can disable both LDP graceful restart and LDP graceful restart helper mode.

## Disabling LDP Graceful Restart or Helper Mode

To disable LDP graceful restart and recovery, include the **disable** statement:

```
ldp {
 graceful-restart {
 disable;
 }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

You can disable helper mode at the LDP protocols level only. You cannot disable helper mode for a specific routing instance. To disable LDP helper mode, include the **helper-disable** statement:

```
ldp {
 graceful-restart {
 helper-disable;
 }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The following LDP graceful restart configurations are possible:

- LDP graceful restart and helper mode are both enabled.
- LDP graceful restart is disabled but helper mode is enabled. A router configured in this way cannot restart gracefully but can help a restarting neighbor.
- LDP graceful restart and helper mode are both disabled. The router does not use LDP graceful restart or the graceful restart type, length, and value (TLV) sent in the initialization message. The router behaves as a router that cannot support LDP graceful restart.

A configuration error is issued if you attempt to enable graceful restart and disable helper mode.

## Configuring Reconnect Time

After the LDP connection between neighbors fails, neighbors wait a certain amount of time for the gracefully restarting router to resume sending LDP messages. After the wait period, the LDP session can be reestablished. You can configure the wait period in seconds. This value is included in the fault tolerant session TLV sent in LDP initialization messages when LDP graceful restart is enabled.

Suppose that Router A and Router B are LDP neighbors. Router A is the restarting Router. The reconnect time is the time that Router A tells Router B to wait after Router B detects that Router A restarted.

To configure the reconnect time, include the **reconnect-time** statement:

```
graceful-restart {
 reconnect-time seconds;
}
```

You can set the reconnect time to a value in the range from 30 through 300 seconds. By default, it is 60 seconds.

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

## Configuring Recovery Time and Maximum Recovery Time

The recovery time is the amount of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This period is also typically the amount of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

To prevent a neighboring router from being adversely affected if it receives a false value for the recovery time from the restarting router, you can configure the maximum recovery time on the neighboring router. A neighboring router maintains its state for the shorter of the two times. For example, Router A is performing an LDP graceful restart. It has sent a recovery time of 900 seconds to neighboring Router B. However, Router B has its maximum recovery time configured at 400 seconds. Router B will only wait for 400 seconds before it purges its LDP information from Router A.

To configure recovery time, include the **recovery-time** statement and the **maximum-neighbor-recovery-time** statement:

```
graceful-restart {
 maximum-neighbor-recovery-time seconds;
 recovery-time seconds;
}
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

## Filtering Inbound LDP Label Bindings

You can filter received LDP label bindings, applying policies to accept or deny bindings advertised by neighboring routers. To configure received-label filtering, include the **import** statement:

```
import [policy-names];
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The named policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings received from all LDP neighbors. All filtering is done with **from** statements. [Table 15 on page 536](#) lists the only **from** operators that apply to LDP received-label filtering.

**Table 15: from Operators That Apply to LDP Received-Label Filtering**

| from Operator       | Description                                                                                |
|---------------------|--------------------------------------------------------------------------------------------|
| <b>interface</b>    | Matches on bindings received from a neighbor that is adjacent over the specified interface |
| <b>neighbor</b>     | Matches on bindings received from the specified LDP router ID                              |
| <b>next-hop</b>     | Matches on bindings received from a neighbor advertising the specified interface address   |
| <b>route-filter</b> | Matches on bindings with the specified prefix                                              |

If a binding is filtered, it still appears in the LDP database, but is not considered for installation as part of a label-switched path (LSP).

Generally, applying policies in LDP can be used only to block the establishment of LSPs, not to control their routing. This is because the path that an LSP follows is determined by unicast routing, and not by LDP. However, when there are multiple equal-cost paths to the destination through different neighbors, you can use LDP filtering to exclude some of the possible next hops from consideration. (Otherwise, LDP chooses one of the possible next hops at random.)

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router (not per-interface) labels; so if multiple parallel links exist between two routers, only one LDP session is established, and it is not bound to a single interface. When a router has multiple adjacencies to the same neighbor, take care to ensure that the filter does what is expected. (Generally, using **next-hop** and **interface** is not appropriate in this case.)

If a label has been filtered (meaning that it has been rejected by the policy and is not used to construct an LSP), it is marked as filtered in the database:

```
user@host> show ldp database
Input label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.6/32 (Filtered)
Output label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.1/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

## Examples: Filtering Inbound LDP Label Bindings

Accept only /32 prefixes from all neighbors:

```
[edit]
protocols {
 ldp {
 import only-32;
 ...
 }
}
policy-options {
 policy-statement only-32 {
 term first {
 from {
 route-filter 0.0.0.0/0 upto /31;
 }
 then reject;
 }
 then accept;
 }
}
```

Accept **131.108/16** or longer from router ID **10.10.255.2** and accept all prefixes from all other neighbors:

```
[edit]
protocols {
 ldp {
 import nosy-neighbor;
 ...
 }
}
policy-options {
 policy-statement nosy-neighbor {
 term first {
 from {
 neighbor 10.10.255.2;
 route-filter 131.108.0.0/16 orlonger accept;
 route-filter 0.0.0.0/0 orlonger reject;
 }
 }
 then accept;
 }
}
```

## Filtering Outbound LDP Label Bindings

You can configure export policies to filter LDP outbound labels. You can filter outbound label bindings by applying routing policies to block bindings from being advertised to neighboring routers. To configure outbound label filtering, include the **export** statement:

**export** [*policy-name*];

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The named export policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings transmitted to all LDP neighbors. The only **from** operator that applies

to LDP outbound label filtering is **route-filter**, which matches bindings with the specified prefix. The only **to** operators that apply to outbound label filtering are the operators in [Table 16 on page 538](#).

**Table 16: to Operators for LDP Outbound-Label Filtering**

| to Operator      | Description                                                                          |
|------------------|--------------------------------------------------------------------------------------|
| <b>interface</b> | Matches on bindings sent to a neighbor that is adjacent over the specified interface |
| <b>neighbor</b>  | Matches on bindings sent to the specified LDP router ID                              |
| <b>next-hop</b>  | Matches on bindings sent to a neighbor advertising the specified interface address   |

If a binding is filtered, the binding is not advertised to the neighboring router, but it can be installed as part of an LSP on the local router. You can apply policies in LDP to block the establishment of LSPs, but not to control their routing. The path an LSP follows is determined by unicast routing, not by LDP.

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router (not per-interface) labels. If multiple parallel links exist between two routers, only one LDP session is established, and it is not bound to a single interface.

Do not use the **next-hop** and **interface** operators when a router has multiple adjacencies to the same neighbor.

Filtered labels are marked in the database:

```
user@host> show ldp database
Input label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
100007 10.10.255.2/32
3 10.10.255.3/32
Output label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
3 10.10.255.1/32
100001 10.10.255.6/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

## Examples: Filtering Outbound LDP Label Bindings

Block transmission of the route for **10.10.255.6/32** to any neighbors:

```
[edit protocols]
ldp {
 export block-one;
}
policy-options {
 policy-statement block-one {
 term first {
 from {
```



```

 route-filter 10.10.255.6/32 exact;
 }
 then reject;
}
then accept;
}
}

```

Send only 131.108/16 or longer to router ID 10.10.255.2, and send all prefixes to all other routers:

```

[edit protocols]
ldp {
 export limit-lsps;
}
policy-options {
 policy-statement limit-lsps {
 term allow-one {
 from {
 route-filter 131.108.0.0/16 orlonger;
 }
 to {
 neighbor 10.10.255.2;
 }
 then accept;
 }
 term block-the-rest {
 to {
 neighbor 10.10.255.2;
 }
 then reject;
 }
 then accept;
 }
}
}

```

## Specifying the Transport Address Used by LDP

Routers must first establish a TCP session between each other before they can establish an LDP session. The TCP session enables the routers to exchange the label advertisements needed for the LDP session. To establish the TCP session, each router must learn the other router's transport address. The transport address is an IP address used to identify the TCP session over which the LDP session will run.

To configure the LDP transport address, include the `transport-address` statement:

**transport-address** (router-id | interface);

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

If you specify the **router-id** option, the address of the router identifier is used as the transport address (unless otherwise configured, the router identifier is typically the same as the loopback address). If you specify the **interface** option, the interface address is used as the transport address for any LDP sessions to neighbors that can be reached over that interface. Note that the router identifier is used as the transport address by default.

You cannot specify the **interface** option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the **router-id** option.

Related Documentation

- [transport-address on page 1083](#)

## Configuring the Prefixes Advertised into LDP from the Routing Table

You can control the set of prefixes that are advertised into LDP and cause the router to be the egress router for those prefixes. By default, only the loopback address is advertised into LDP. To configure the set of prefixes from the routing table to be advertised into LDP, include the **egress-policy** statement:

```
egress-policy policy-name;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.



**NOTE:** If you configure an egress policy for LDP that does not include the loopback address, it is no longer advertised in LDP. To continue to advertise the loopback address, you need to explicitly configure it as a part of the LDP egress policy.

The named policy (configured at the **[edit policy-options]** or **[edit logical-systems logical-system-name policy-options]** hierarchy level) is applied to all routes in the routing table. Those routes that match the policy are advertised into LDP. You can control the set of neighbors to which those prefixes are advertised by using the **export** statement. Only **from** operators are considered; you can use any valid **from** operator. For more information, see the *Junos OS Routing Protocols Library for Routing Devices*.

### Example: Configuring the Prefixes Advertised into LDP

Advertise all connected routes into LDP:

```
[edit protocols]
ldp {
 egress-policy connected-only;
}
policy-options {
 policy-statement connected-only {
 from {
 protocol direct;
 }
 then accept;
 }
}
```

## Configuring FEC Deaggregation

When an LDP egress router advertises multiple prefixes, the prefixes are bound to a single label and aggregated into a single forwarding equivalence class (FEC). By default, LDP maintains this aggregation as the advertisement traverses the network.

Normally, because an LSP is not split across multiple next hops and the prefixes are bound into a single LSP, load-balancing across equal-cost paths does not occur. You can, however, load-balance across equal-cost paths if you configure a load-balancing policy and deaggregate the FECs.

Deaggregating the FECs causes each prefix to be bound to a separate label and become a separate LSP.

To configure deaggregated FECs, include the **deaggregate** statement:

```
deaggregate;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For all LDP sessions, you can configure deaggregated FECs only globally.

Deaggregating a FEC allows the resulting multiple LSPs to be distributed across multiple equal-cost paths and distributes LSPs across the multiple next hops on the egress segments but installs only one next hop per LSP.

To aggregate FECs, include the **no-deaggregate** statement:

```
no-deaggregate;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

For all LDP sessions, you can configure aggregated FECs only globally.

### Related Documentation

- [Configuring Load Balancing Across RSVP LSPs on page 481](#)
- [Configuring Protocol-Independent Load Balancing in Layer 3 VPNs](#)
- [Configuring VPLS Load Balancing](#)
- [Example: Load Balancing BGP Traffic](#)

## Configuring Policers for LDP FECs

You can configure the Junos OS to track and police traffic for LDP FECs. LDP FEC policers can be used to do any of the following:

- Track or police the ingress traffic for an LDP FEC.
- Track or police the transit traffic for an LDP FEC.
- Track or police LDP FEC traffic originating from a specific forwarding class.

- Track or police LDP FEC traffic originating from a specific virtual routing and forwarding (VRF) site.
- Discard false traffic bound for a specific LDP FEC.

To police traffic for an LDP FEC, you must first configure a filter. Specifically, you need to configure either the **interface** statement or the **interface-set** statement at the **[edit firewall family protocol-family filter filter-name term term-name from]** hierarchy level. The **interface** statement allows you to match the filter to a single interface. The **interface-set** statement allows you to match the filter to multiple interfaces.

For more information on how to configure the **interface** statement, the **interface-set** statement, and policers for LDP FECs, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

Once you have configured the filters, you need to include them in the **policing** statement configuration for LDP. To configure policers for LDP FECs, include the **policing** statement:

```
policing {
 fec fec-address {
 ingress-traffic filter-name;
 transit-traffic filter-name;
 }
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

The **policing** statement includes the following options:

- **fec**—Specify the FEC address for the LDP FEC you want to police.
- **ingress-filter**—Specify the name of the ingress traffic filter.
- **transit-traffic**—Specify the name of the transit traffic filter.

---

## Configuring LDP IPv4 FEC Filtering

By default, when a targeted LDP session is established, the Junos OS always exchanges both the IPv4 forwarding equivalence classes (FECs) and the Layer 2 circuit FECs over the targeted LDP session. For an LDP session to an indirectly connected neighbor, you might only want to export Layer 2 circuit FECs to the neighbor if the session was specifically configured to support Layer 2 circuits or VPLS.

In a mixed vendor network where all non-BGP prefixes are advertised into LDP, the LDP database can become large. For this type of environment, it can be useful to prevent the advertisement of IPv4 FECs over LDP sessions formed because of Layer 2 circuit or LDP VPLS configuration. Similarly, it can be useful to filter any IPv4 FECs received in this sort of environment.

If all the LDP neighbors associated with an LDP session are Layer 2 only, you can configure the Junos OS to advertise only Layer 2 circuit FECs by configuring the **l2-smart-policy**

statement. This feature also automatically filters out the IPv4 FECs received on this session. If you have configured an explicit export or import policy, this feature is disabled.

If one of the LDP session's neighbors is formed because of a discovered adjacency or if the adjacency is formed because of an LDP tunneling configuration on one or more RSVP LSPs, the IPv4 FECs are advertised and received using the default behavior.

To prevent LDP from exporting IPv4 FECs over LDP sessions with Layer 2 neighbors only and to filter out IPv4 FECs received over such sessions, include the **l2-smart-policy** statement:

```
l2-smart-policy;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary for this statement.

## Configuring BFD for LDP LSPs

You can configure Bidirectional Forwarding Detection (BFD) for LDP LSPs. The BFD protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the router stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. The failure detection timers for BFD have shorter time limits than the failure detection mechanisms of static routes, providing faster detection.

An error is logged whenever a BFD session for a path fails. The following shows how BFD for LDP LSP log messages might appear:

```
RPD_LDP_BFD_UP: LDP BFD session for FEC 10.255.16.14/32 is up
RPD_LDP_BFD_DOWN: LDP BFD session for FEC 10.255.16.14/32 is down
```

You can also configure BFD for RSVP LSPs, as described in [“Configuring BFD for RSVP-Signaled LSPs” on page 355](#).

The BFD failure detection timers are adaptive and can be adjusted to be more or less aggressive. For example, the timers can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

To enable BFD for LDP LSPs, include the **oam** and **bfd-liveness-detection** statements:

```
oam {
 bfd-liveness-detection {
 detection-time threshold milliseconds;
 ecmp;
 failure-action {
 remove-nexthop;
 remove-route;
```

```

 }
 holddown-interval seconds;
 ingress-policy ingress-policy-name;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 minimum-transmit-interval milliseconds;
 multiplier detection-time-multiplier;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (0 | 1 | automatic);
}
fec fec-address {
 bfd-liveness-detection {
 detection-time threshold milliseconds;
 ecmp;
 failure-action {
 remove-nexthop;
 remove-route;
 }
 holddown-interval milliseconds;
 ingress-policy ingress-policy-name;
 minimum-interval milliseconds;
 minimum-receive-interval milliseconds;
 minimum-transmit-interval milliseconds;
 multiplier detection-time-multiplier;
 no-adaptation;
 transmit-interval {
 minimum-interval milliseconds;
 threshold milliseconds;
 }
 version (0 | 1 | automatic);
 }
 no-bfd-liveness-detection;
 periodic-traceroute {
 disable;
 exp exp-value;
 fanout fanout-value;
 frequency minutes;
 paths number-of-paths;
 retries retry-attempts;
 source address;
 ttl tvl-value;
 wait seconds;
 }
}
lsp-ping-interval seconds;
periodic-traceroute {
 disable;
 exp exp-value;
 fanout fanout-value;
 frequency minutes;
 paths number-of-paths;
 retries retry-attempts;
}

```

```

 source address;
 ttl ttl-value;
 wait seconds;
 }
}

```

You can enable BFD for the LDP LSPs associated with a specific forwarding equivalence class (FEC) by configuring the FEC address using the **fec** option at the **[edit protocols ldp]** hierarchy level. Alternatively, you can configure an Operation Administration and Management (OAM) ingress policy to enable BFD on a range of FEC addresses. For more information, see [“Configuring OAM Ingress Policies for LDP” on page 547](#).

You cannot enable BFD LDP LSPs unless their equivalent FEC addresses are explicitly configured or OAM is enabled on the FECs using an OAM ingress policy. If BFD is not enabled for any FEC addresses, the BFD session will not come up.

You can configure the **oam** statement at the following hierarchy levels:

- **[edit protocols ldp]**
- **[edit logical-systems *logical-system-name* protocols ldp]**

The **oam** statement includes the following options:

- **fec**—Specify the FEC address. You must either specify a FEC address or configure an OAM ingress policy to ensure that the BFD session comes up.
- **lsp-ping-interval**—Specify the duration of the LSP ping interval in seconds. To issue a ping on an LDP-signaled LSP, use the **ping mpls ldp** command. For more information, see the [CLI Explorer](#).

The **bfd-liveness-detection** statement includes the following options:

- **ecmp**—Cause LDP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the **ecmp** option, you must also configure the **periodic-traceroute** statement for the specified FEC. If you do not do so, the commit operation fails. You can configure the **periodic-traceroute** statement at the global hierarchy level (**[edit protocols ldp oam]**) while only configuring the **ecmp** option for a specific FEC (**[edit protocols ldp oam fec address bfd-liveness-detection]**).
- **holddown-interval**—Specify the duration the BFD session should remain up before adding the route or next hop. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.
- **minimum-interval**—Specify the minimum transmit and receive interval. If you configure the **minimum-interval** option, you do not need to configure the **minimum-receive-interval** option or the **minimum-transmit-interval** option.
- **minimum-receive-interval**—Specify the minimum receive interval. The range is from 1 through 255,000 milliseconds.
- **minimum-transmit-interval**—Specify the minimum transmit interval. The range is from 1 through 255,000 milliseconds.

- **multiplier**—Specify the detection time multiplier. The range is from 1 through 255.
- **version**—Specify the BFD version. The options are BFD version 0 or BFD version 1. By default, the Junos OS software attempts to automatically determine the BFD version.

## Configuring ECMP-Aware BFD for LDP LSPs

---

When you configure BFD for a FEC, a BFD session is established for only one active local next-hop for the router. However, you can configure multiple BFD sessions, one for each FEC associated with a specific equal-cost multipath (ECMP) path. For this to function properly, you also need to configure LDP LSP periodic traceroute. (See [“Configuring LDP LSP Traceroute” on page 634](#).) LDP LSP traceroute is used to discover ECMP paths. A BFD session is initiated for each ECMP path discovered. Whenever a BFD session for one of the ECMP paths fails, an error is logged.

LDP LSP traceroute is run periodically to check the integrity of the ECMP paths. The following might occur when a problem is discovered:

- If the latest LDP LSP traceroute for a FEC differs from the previous traceroute, the BFD sessions associated with that FEC (the BFD sessions for address ranges that have changed from previous run) are brought down and new BFD sessions are initiated for the destination addresses in the altered ranges.
- If the LDP LSP traceroute returns an error (for example, a timeout), all the BFD sessions associated with that FEC are torn down.

To configure LDP to establish BFD sessions for all ECMP paths configured for the specified FEC, include the **ecmp** statement.

```
ecmp;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Along with the **ecmp** statement, you must also include the **periodic-traceroute** statement, either in the global LDP OAM configuration (at the **[edit protocols ldp oam]** or **[edit logical-systems logical-system-name protocols ldp oam]** hierarchy level) or in the configuration for the specified FEC (at the **[edit protocols ldp oam fec address]** or **[edit logical-systems logical-system-name protocols ldp oam fec address]** hierarchy level). Otherwise, the commit operation fails.

## Configuring a Failure Action for the BFD Session on an LDP LSP

---

You can configure route and next-hop properties in the event of a BFD session failure event on an LDP LSP. The failure event could be an existing BFD session that has gone down or could be a BFD session that never came up. LDP adds back the route or next hop when the relevant BFD session comes back up.



You can configure one of the following failure action options for the **failure-action** statement in the event of a BFD session failure on the LDP LSP:

- **remove-nexthop**—Removes the route corresponding to the next hop of the LSP's route at the ingress node when a BFD session failure event is detected.
- **remove-route**—Removes the route corresponding to the LSP from the appropriate routing tables when a BFD session failure event is detected. If the LSP is configured with ECMP and a BFD session corresponding to any path goes down, the route is removed.

To configure a failure action in the event of a BFD session failure on an LDP LSP, include either the **remove-nexthop** option or the **remove-route** option for the **failure-action** statement:

```
failure-action {
 remove-nexthop;
 remove-route;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Configuring the Holddown Interval for the BFD Session

You can specify the duration the BFD session should be up before adding a route or next hop by configuring the **holddown-interval** statement at either the **[edit protocols ldp oam bfd-liveness-detection]** hierarchy level or at the **[edit protocols ldp oam fec address bfd-liveness-detection]** hierarchy level. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.

```
holddown-interval seconds;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Configuring OAM Ingress Policies for LDP

Using the **ingress-policy** statement, you can configure an Operation, Administration, and Management (OAM) policy to choose which forwarding equivalence classes (FECs) need to have OAM enabled. If the FEC passes through the policy or if the FEC is explicitly configured, OAM is enabled for a FEC. For FECs chosen using a policy, the BFD parameters configured under **[edit protocols ldp oam bfd-liveness-detection]** are applied.

You configure the OAM ingress policy at the **[edit policy-options]** hierarchy level. To configure an OAM ingress policy, include the **ingress-policy** statement:

```
ingress-policy ingress-policy-name;
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols ldp oam]**
- **[edit logical-systems logical-system-name protocols ldp oam]**

## Configuring LDP Link Protection

---

You can configure Label Distribution Protocol (LDP) link protection for both unicast and multicast LDP label-switched paths (LSPs) to provide resiliency during link or node failure.

Before you begin:

1. Configure the device interfaces.
2. Configure the router ID and autonomous system number for the device.
3. Configure the following protocols:
  - a. RSVP
  - b. MPLS with traffic engineering capability.
  - c. OSPF with traffic engineering capability.



**NOTE:** For multicast LDP link protection with loop-free alternative (LFA), enable link protection.

[edit protocols]

user@R0# set ospf area 0 interface all link-protection

---

To configure LDP link protection:

1. Enable point-to-multipoint LDP LSPs.

[edit protocols]

user@R0# set ldp p2mp

2. Enable LDP on all the interfaces of Router R0 (excluding the management interface) and configure link protection with dynamic RSVP bypass LSP.

[edit protocols]

user@R0# set ldp interface all link-protection dynamic-rsvp-lsp

user@R0# set ldp interface fxp0.0 disable

3. Verify and commit the configuration.

For example:

[edit protocols]

user@R0# show protocols

rsvp {

  interface all;

  interface fxp0.0 {

    disable;

  }

}

mpls {

  traffic-engineering;

  interface all;

  interface fxp0.0 {

    disable;

```

 }
 }
 ospf {
 traffic-engineering;
 area 0.0.0.0 {
 interface all {
 metric 1;
 }
 interface fxp0.0 {
 disable;
 }
 }
 }
 ldp {
 interface all {
 link-protection {
 dynamic-rsvp-lsp;
 }
 }
 interface fxp0.0 {
 disable;
 }
 p2mp;
 }

[edit]
user@R0# commit
commit complete

```

**Related Documentation**

- [Example: Configuring LDP Link Protection on page 549](#)

## Example: Configuring LDP Link Protection

- [LDP Link Protection Overview on page 549](#)

### LDP Link Protection Overview

- [Introduction to LDP on page 550](#)
- [Junos OS LDP Protocol Implementation on page 550](#)
- [Understanding Multipoint Extensions to LDP on page 550](#)
- [Using Multipoint Extensions to LDP on Targeted LDP Sessions on page 551](#)
- [Current Limitations of LDP Link Protection on page 552](#)
- [Using RSVP LSP as a Solution on page 553](#)
- [Understanding Multicast LDP Link Protection on page 555](#)
- [Different Modes for Providing LDP Link Protection on page 555](#)
- [Label Operation for LDP Link Protection on page 557](#)
- [Sample Multicast LDP Link Protection Configuration on page 563](#)

- [Make-Before-Break on page 564](#)
- [Caveats and Limitations on page 566](#)

---

## Introduction to LDP

The Label Distribution Protocol (LDP) is a protocol for distributing 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 the data link LSPs.

These LSPs might have an endpoint at a directly attached neighbor (comparable to IP hop-by-hop forwarding) or at a network egress node, enabling switching through all intermediary nodes. LSPs established by LDP can also traverse traffic-engineered LSPs created by RSVP.

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers. This process forms a tree of LSPs that converge on the egress router.

---

## Junos OS LDP Protocol Implementation

The Junos OS implementation of LDP supports LDP version 1. Junos OS supports a simple mechanism for tunneling between routers in an interior gateway protocol (IGP), to eliminate the required distribution of external routes within the core. Junos OS allows an MPLS tunnel next hop to all egress routers in the network, with only an IGP running in the core to distribute routes to egress routers. Edge routers run BGP but do not distribute external routes to the core. Instead, the recursive route lookup at the edge resolves to an LSP switched to the egress router. No external routes are necessary on the transit LDP routers.

---

## Understanding Multipoint Extensions to LDP

An LDP defines mechanisms for setting up point-to-point, multipoint-to-point, point-to-multipoint, and multipoint-to-multipoint LSPs in the network. The point-to-multipoint and multipoint-to-multipoint LSPs are collectively referred to as multipoint LSPs, where traffic flows from a single source to multiple destinations, and from multiple sources to multiple destinations, respectively. The destination or egress routers are called leaf nodes, and traffic from the source traverses one or more transit nodes before reaching the leaf nodes.



**NOTE:** Junos OS does not provide support for multipoint-to-multipoint LSPs.

---


By taking advantage of the MPLS packet replication capability of the network, multipoint LSPs avoid unnecessary packet replication at the ingress router. Packet replication takes place only when packets are forwarded to two or more different destinations requiring different network paths.

### Using Multipoint Extensions to LDP on Targeted LDP Sessions

The specification for the multipoint extensions to LDP requires that the two endpoints of an LDP session are directly connected by a Layer 2 medium, or are considered to be neighbors by the network's IGP. This is referred to as an LDP link session. When the two endpoints of an LDP session are not directly connected, the session is referred to as a targeted LDP session.

Past Junos OS implementations support multicast LDP for link sessions only. With the introduction of the LDP link protection feature, the multicast LDP capabilities are extended to targeted LDP sessions. [Figure 48 on page 551](#) shows a sample topology.

**Figure 48: Multicast LDP Support for Targeted LDP Session**



Routers R7 and R8 are the upstream (LSR-U) and downstream (LSR-D) label-switched routers (LSRs), respectively, and deploy multicast LDP. The core router, Router R5, has RSVP-TE enabled.

When LSR-D is setting up the point-to-multipoint LSP with root and LSP ID attributes, it determines the upstream LSR-U as a next-hop on the best path to the root (currently, this next-hop is assumed to be an IGP next hop).

With the multicast LDP support on targeted LDP sessions, you can determine if there is an LSP next hop to LSR-U which is on LSR-D's path to root, where LSR-D and LSR-U are not directly connected neighbors, but targeted LDP peers. The point-to-multipoint label advertised on the targeted LDP session between LSR-D and LSR-U is not used unless there is an LSP between LSR-D and LSR-U. Therefore, a corresponding LSP in the reverse direction from LSR-U to LSR-D is required.

Data is transmitted on the point-to-multipoint LSP using unicast replication of packets, where LSR-U sends one copy to each downstream LSR of the point-to-multipoint LSP.

The data transmission is implemented in the following ways:

1. The point-to-multipoint capabilities on the targeted LDP session are negotiated.
2. The algorithm to select the upstream LSR is changed, where if IGP next hops are unavailable, or in other words, there is no LDP link session between LSR-D and LSR-U, an RSVP LSP is used as the next hop to reach LSR-U.

3. The incoming labels received over the targeted LDP sessions are installed as a branch next hop for this point-to-multipoint FEC route with the LDP label as the inner label and the RSVP label as the outer label.

### Current Limitations of LDP Link Protection

When there is a link or node failure in an LDP network deployment, fast traffic recovery should be provided to recover impacted traffic flows for mission-critical services. In the case of multipoint LSPs, when one of the links of the point-to-multipoint tree fails, the subtrees might get detached until the IGP reconverges and the multipoint LSP is established using the best path from the downstream router to the new upstream router.


In fast reroute using local repair for LDP traffic, a backup path (repair path) is pre-installed in the Packet Forwarding Engine. When the primary path fails, traffic is rapidly moved to the backup path without having to wait for the routing protocols to converge. Loop-free alternate (LFA) is one of the methods used to provide IP fast reroute capability in the core and service provider networks.

Without LFA, when a link or a router fails or is returned to service, the distributed routing algorithms compute the new routes based on the changes in the network. The time during which the new routes are computed is referred to as routing transition. Until the routing transition is completed, the network connectivity is interrupted because the routers adjacent to a failure continue to forward the data packets through the failed component until an alternative path is identified.

However, LFA does not provide full coverage in all network deployments because of the IGP metrics. As a result, this is a limitation to the current LDP link protection schemes.

Figure 49 on page 552 illustrates a sample network with incomplete LFA coverage, where traffic flows from the source router (S) to the destination router (D) through Router R1. Assuming that each link in the network has the same metric, if the link between the Router S and Router R1 fails, Router R4 is not an LFA that protects the S-R1 link, so traffic resiliency is lost. Thus, full coverage is not achieved by using plain LFA. In typical networks, there is always some percentage of LFA coverage gap with plain LFA.

Figure 49: Incomplete Coverage Problem with LFA



### Using RSVP LSP as a Solution

The key to protect the traffic flowing through LDP LSPs is to have an explicit tunnel to re-route the traffic in the event of a link or node failure. The explicit path has to terminate on the next downstream router, and the traffic needs to be accepted on the explicit path, where the RPF check should pass.


RSVP LSPs help overcome the current limitations of loop-free alternate (LFA) for both point-to-point and point-to-multipoint LDP LSPs by extending the LFA coverage in the following methods:

- [Manually Configured RSVP LSPs on page 553](#)
- [Dynamically Configured RSVP LSPs on page 553](#)

#### Manually Configured RSVP LSPs

Considering the example used in [Figure 49 on page 552](#), when the S-R1 link fails, and Router R4 is not an LFA for that particular link, a manually created RSVP LSP is used as a patch to provide complete LFA coverage. The RSVP LSP is pre-sigaled and pre-installed in the Packet Forwarding Engine of Router S, so that it can be used as soon as Router S detects that the link has failed.

Figure 50: Manually Configured RSVP LSP Coverage



In this case, an RSVP LSP is created between Routers S, R4, and R3 as illustrated in [Figure 50 on page 553](#). A targeted LDP session is created between Router S and Router R3, as a result of which, when the S-R1 link fails, traffic reaches Router R3. Router R3 forwards the traffic to Router R2, as it is the shortest path to reach the destination, Router D.


#### Dynamically Configured RSVP LSPs

In this method, the RSVP LSPs are created automatically and pre-installed in the system so that they can be used immediately when there is a link failure. Here, the egress is the node on the other side of the link being protected, thereby improving the LFA coverage.

Considering the example used in [Figure 49 on page 552](#), in order to protect traffic against the potential failure of the S-R1 link, because Router R4 is not an LFA for that particular link, an RSVP bypass LSP is automatically created to Router R1, which is the node on the far side of the protected link as illustrated in [Figure 51 on page 554](#). From Router R1, traffic is forwarded to its original destination, Router D.

The RSVP LSP is pre-sigaled and pre-installed in the Packet Forwarding Engine of Router S so that it can be used as soon as Router S detects that the link has failed.

**Figure 51: Dynamically Configured RSVP LSP Coverage**



An alternative mode of operation is not to use LFA at all, and to always have the RSVP LSP created to cover all link failures.

To enable dynamic RSVP LSPs, include the **dynamic-rsvp-lsp** statement at the **[edit protocols ldp interface *interface-name* link-protection]** hierarchy level, in addition to enabling the RSVP protocol on the appropriate interfaces.

Some of the benefits of enabling dynamic RSVP LSPs include:

- Ease of configuration.
- 100 percent coverage against link failure as long as there is an alternate path to the far end of the link being protected.
- Setting up and tearing down of the RSVP bypass LSP is automatic.
- RSVP LSP only used for link protection and not for forwarding traffic while the link being protected is up.
- Reduces the total number of RSVP LSPs required on the system.



### Understanding Multicast LDP Link Protection

A point-to-multipoint LDP label-switched path (LSP) is an LDP-signaled LSP that is point-to-multipoint, and is referred to as multicast LDP.

A multicast LDP LSP can be used to send traffic from a single root or ingress node to a number of leaf or egress nodes traversing one or more transit nodes. Multicast LDP link protection enables fast reroute of traffic carried over point-to-multipoint LDP LSPs in case of a link failure. When one of the links of the point-to-multipoint tree fails, the subtrees might get detached until the IGP reconverges and the multipoint LSP is established using the best path from the downstream router to the new upstream router.

To protect the traffic flowing through the multicast LDP LSP, you can configure an explicit tunnel to re-route the traffic in the event of link failure. The explicit path has to terminate on the next downstream router. The reverse path forwarding for the traffic should be successful.

Multicast LDP link protection introduces the following features and functionality:

- Use of dynamic RSVP LSP as bypass tunnels

The RSVP LSP's Explicit Route Object (ERO) is calculated using Constrained Shortest Path First (CSPF) with the constraint as the link to avoid. The LSP is signaled and torn down dynamically whenever link protection is necessary.

- Make-before-break

The make-before-break feature ensures that there is minimum packet loss when attempting to signal a new LSP path before tearing down the old LSP path for the multicast LDP LSP.

- Targeted LDP session

A targeted adjacency to the downstream label-switching router (LSR) is created for two reasons:

- To keep the session up after link failure.
- To use the point-to-multipoint label received from the session to send traffic to the downstream LSR on the RSVP LSP bypass tunnel.

When the downstream LSR sets up the multicast LDP LSP with the root node and LSP ID, it uses that upstream LSR, which is on the best path toward the root.



**NOTE:** Multicast LDP link protection is not required when there are multiple link adjacencies (parallel links) to the downstream LSR.

### Different Modes for Providing LDP Link Protection

Following are three different modes of operation available for unicast and multicast LDP link protection:

- **Case A: LFA only**

Under this mode of operation, multicast LDP link protection is provided using an existing viable loop-free alternate (LFA). In the absence of a viable LFA, link protection is not provided for the multicast LDP LSP.

- **Case B: LFA and Dynamic RSVP LSP**


Under this mode of operation, multicast LDP link protection is provided using an existing viable LFA. In the absence of a viable LFA, an RSVP bypass LSP is created automatically to provide link protection for the multicast LDP LSP.

- **Case C: Dynamic RSVP LSP only**

Under this mode of operation, LFA is not used for link protection. Multicast LDP link protection is provided by using automatically created RSVP bypass LSP.

Figure 52 on page 556 is a sample topology illustrating the different modes of operation for multicast LDP link protection. Router R5 is the root connecting to two leaf nodes, Routers R3 and R4. Router R0 and Router R1 are the upstream and downstream label-switched routers (LSRs), respectively. A multicast LDP LSP runs among the root and leaf nodes.

**Figure 52: Multicast LDP Link Protection Sample Topology**



Considering that Router R0 needs to protect the multicast LDP LSP in the case that the R0-R1 link fails, the different modes of link protection operate in the following manner:

- **Case A: LFA only**

Router R0 checks if a viable LFA path exists that can avoid the R0-R1 link to reach Router R1. Based on the metrics, Router R2 is a valid LFA path for the R0-R1 link and is used to forward unicast LDP traffic. If multiple multicast LDP LSPs use the R0-R1 link, the same LFA (Router R2) is used for multicast LDP link protection.

When the R0-R1 link fails, the multicast LDP LSP traffic is moved onto the LFA path by Router R0, and the unicast LDP label to reach Router R1 (L100) is pushed on top of the multicast LDP label (L21).

- **Case B: LFA and Dynamic RSVP LSP**

Router R0 checks if a viable LFA path exists that can avoid the R0-R1 link to reach Router R1. Based on the metrics, Router R2 is a valid LFA path for the R0-R1 link and is used to forward unicast LDP traffic. If multiple multicast LDP LSPs use the R0-R1 link, the same LFA (Router R2) is used for multicast LDP link protection. When the R0-R1 link fails, the multicast LDP LSP traffic is moved onto the LFA path by Router R0.

However, if the metric on the R2-R1 link was 50 instead of 10, Router 2 is not a valid LFA for the R0-R1 link. In this case, an RSVP LSP is automatically created to protect the multicast LDP traffic traveling between Routers R0 and R1.

- **Case C: Dynamic RSVP LSP only**


An RSVP LSP is signaled automatically from Router R0 to Router R1 through Router R2, avoiding interface ge-1/1/0. If multiple multicast LDP LSPs use the R0-R1 link, the same RSVP LSP is used for multicast LDP link protection.

When the R0-R1 link fails, the multicast LDP LSP traffic is moved onto the RSVP LSP by Router R0, and the RSVP label to reach Router R1 (L100) is pushed on top of the multicast LDP label (L21).

### Label Operation for LDP Link Protection

Using the same network topology as in Figure 5, [Figure 53 on page 557](#) illustrates the label operation for unicast and multicast LDP link protection.

**Figure 53: LDP Label Operation Sample Topology**



Router R5 is the root connecting to two leaf nodes, Routers R3 and R4. Router R0 and Router R1 are the upstream and downstream label-switched routers (LSRs), respectively. A multicast LDP LSP runs among the root and leaf nodes. An unicast LDP path connects Router R1 to Router R5.

The label operation is performed differently under the following modes of LDP link protection:

- [Case A: LFA Only on page 558](#)
- [Case B: LFA and Dynamic RSVP LSP on page 561](#)
- [Case C: Dynamic RSVP LSP Only on page 563](#)

#### **Case A: LFA Only**

Using the **show route detail** command output on Router R0, the unicast LDP traffic and multicast LDP traffic can be derived.

```
user@R0> show route detail
299840 (1 entry, 1 announced)
 *LDP Preference: 9
 Next hop type: Router
 Address: 0x93bc22c
 Next-hop reference count: 1
 Next hop: 11.0.0.6 via ge-0/0/1.0 weight 0x1, selected
 Label operation: Swap 299824
 Session Id: 0x1
 Next hop: 11.0.0.10 via ge-0/0/2.0 weight 0xf000
 Label operation: Swap 299808
 Session Id: 0x3
 State: <Active Int>
 Age: 3:16 Metric:1
 Validation State: unverified
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 Prefixes bound to route: 192.168.0.4/32

299856 (1 entry, 1 announced)
 *LDP Preference: 9
 Next hop type: Flood
 Address: 0x9340e04
 Next-hop reference count: 3
 Next hop type: Router, Next hop index: 262143
 Address: 0x93bc3dc
 Next-hop reference count: 2
 Next hop: 11.0.0.6 via ge-0/0/1.0 weight 0x1
 Label operation: Swap 299888
 Next hop: 11.0.0.10 via ge-0/0/2.0 weight 0xf000
 Label operation: Swap 299888, Push 299776(top)
 Label TTL action: prop-ttl, prop-ttl(top)
 State: <Active Int AckRequest>
 Age: 3:16 Metric:1
 Validation State: unverified
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 FECs bound to route: P2MP root-addr 192.168.0.5, lsp-id 99
```

Label 299840 is traffic arriving at Router R0 that corresponds to unicast LDP traffic to Router R1. Label 299856 is traffic arriving at Router 0 that corresponds to multicast LDP traffic from the root node R5 to the leaf egress nodes, R3 and R4.

The main path for both unicast and multicast LDP LSPs is through interface ge-0/0/1 (the link to Router R1), and the LFA path is through interface ge-0/0/2 (the link to Router R2). The LFA path is not used unless the ge-0/0/1 interface goes down.


In the label operation for Case A, the LFA-only mode of operation is different for unicast and multicast LDP traffic:

- Unicast label operation

For unicast LDP traffic, the FECs and associated labels are advertised on all the links in the network on which LDP is enabled. This means that in order to provide LFA action for the unicast LDP traffic to Router R4, instead of swapping the incoming label for label 299824 advertised by Router R1 for FEC R4, Router R0 simply swaps the incoming label for label 299808 advertised by Router R2 for FEC R4. This is the standard Junos OS LFA operation for unicast LDP traffic.

Figure 54 on page 559 illustrates the label operation for unicast traffic when the R0-R1 link fails. The grey boxes show the label operation for unicast LDP traffic under normal condition, and the dotted boxes show the label operation for unicast LDP traffic when the R0-R1 link fails.

**Figure 54: Unicast LDP Label Operation**



- Multicast label operation

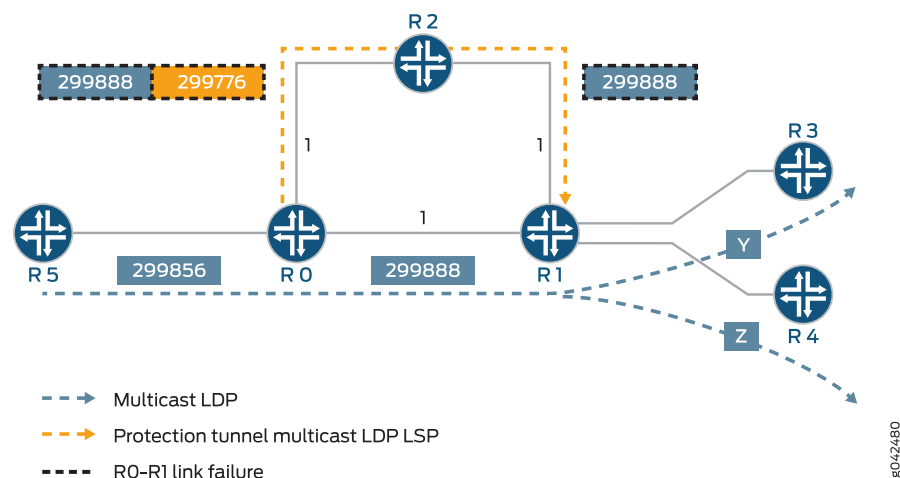
The label operation for multicast LDP traffic differs from the unicast LDP label operation, because multipoint LSP labels are only advertised along the best path from the leaf node to the ingress node. As a result, Router R2 has no knowledge of the multicast LDP. To overcome this, the multicast LDP LSP traffic is simply tunneled inside the unicast LDP LSP path through Router R2 that terminates at Router R1.

In order to achieve this, Router R0 first swaps the incoming multicast LDP LSP label 299856 to label 299888 advertised by Router R1. Label 299776 is then pushed on top, which is the LDP label advertised by Router R2 for FEC R1. When the packet arrives at Router R2, the top label is popped out due to penultimate hop-popping. This means

that the packet arrives at Router R1 with the multicast LDP label 299888 that Router R1 had originally advertised to Router R0.

Figure 55 on page 560 illustrates the label operation for multicast LDP traffic when the R0-R1 link fails. The blue boxes show the label operation for multicast LDP traffic under normal condition, and the dotted boxes show the label operation for multicast LDP traffic when the R0-R1 link fails.

Figure 55: Multicast LDP Label Operation



When the metric on the R2-R1 link is set to 1000 instead of 1, Router R2 is not a valid LFA for Router R0. In this case, if Router R2 receives a packet destined for Router R1, R3, or R4 before its IGP has converged, the packet is sent back to Router R0, resulting in looping packets.

Because Router R0 has no viable LFA, no backup paths are installed in the Packet Forwarding Engine. If the R0-R1 link fails, traffic flow is interrupted until the IGP and LDP converge and new entries are installed on the affected routers.

The **show route detail** command displays the state when no LFA is available for link protection.

```
user@host> show route detail
299840 (1 entry, 1 announced)
 *LDP Preference: 9
 Next hop type: Router, Next hop index: 578
 Address: 0x9340d20
 Next-hop reference count: 2
 Next hop: 11.0.0.6 via ge-0/0/1.0, selected
 Label operation: Swap 299824
 Session Id: 0x1
 State: <Active Int>
 Age: 5:38 Metric: 1
 Validation State: unverified
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 Prefixes bound to route: 192.168.0.4/32
```

```

299856 (1 entry, 1 announced)
 *LDP Preference: 9
 Next hop type: Flood
 Address: 0x9340e04
 Next-hop reference count: 3
 Next hop type: Router, Next hop index: 579
 Address: 0x93407c8
 Next-hop reference count: 2
 Next hop: 11.0.0.6 via ge-0/0/1.0
 Label operation: Swap 299888
 State: <Active Int AckRequest>
 Age: 5:38 Metric: 1
 Validation State: unverified
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 FECs bound to route: P2MP root-addr 192.168.0.5, lsp-id 99

```

### ***Case B: LFA and Dynamic RSVP LSP***

In this mode of operation, if there is a viable LFA neighbor, the label operation behavior is similar to that of Case A, LFA only mode. However, if there is no viable LFA neighbor, an RSVP bypass tunnel is automatically created.

If the metric on the link R2-R1 is set to 1000 instead of 1, Router R2 is not an LFA for Router R0. On learning that there are no LFA paths to protect the R0-R1 link failure, an RSVP bypass tunnel is automatically created with Router R1 as the egress node and follows a path that avoids the R0-R1 link (for instance, R0-R2-R1).

If the R0-R1 link fails, the unicast LDP and multicast LDP traffic is tunneled through the RSVP bypass tunnel. The RSVP bypass tunnel is not used for normal forwarding and is used only to provide link protection to LDP traffic in the case of R0-R1 link failure.

Using the **show route detail** command, the unicast and multicast LDP traffic can be derived.

```

user@host> show route detail
299840 (1 entry, 1 announced)
 *LDP Preference: 9
 Next hop type: Router
 Address: 0x940c3dc
 Next-hop reference count: 1
 Next hop: 11.0.0.6 via ge-0/0/1.0 weight 0x1, selected
 Label operation: Swap 299824
 Session Id: 0x1
 Next hop: 11.0.0.10 via ge-0/0/2.0 weight 0x8001
 Label-switched-path ge-0/0/1.0:BypassLSP->192.168.0.1
 Label operation: Swap 299824, Push 299872(top)
 Label TTL action: prop-ttl, prop-ttl(top)
 Session Id: 0x3
 State: <Active Int NhAckRequest>
 Age: 19 Metric: 1
 Validation State: unverified
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 Prefixes bound to route: 192.168.0.4/32

```

```

299856 (1 entry, 1 announced)
 *LDP Preference: 9
 Next hop type: Flood
 Address: 0x9340e04
 Next-hop reference count: 3
 Next hop type: Router, Next hop index: 262143
 Address: 0x940c154
 Next-hop reference count: 2
 Next hop: 11.0.0.6 via ge-0/0/1.0 weight 0x1
 Label operation: Swap 299888
 Next hop: 11.0.0.10 via ge-0/0/2.0 weight 0x8001
 Label-switched-path ge-0/0/1.0:BypassLSP->192.168.0.1
 Label operation: Swap 299888, Push 299872(top)
 Label TTL action: prop-ttl, prop-ttl(top)
 State: < Active Int AckRequest>
 Age: 20 Metric: 1
 Validation State: unverified
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 FECs bound to route: P2MProot-addr 192.168.0.5, lsp-id 99

```

The main path for both unicast and multicast LDP LSP is through interface ge-0/0/1 (the link to Router R1), and the LFA path is through interface ge-0/0/2 (the link to Router R2). The LFA path is not used unless the ge-0/0/1 interface goes down.


Label 299840 is traffic arriving at Router R0 that corresponds to unicast LDP traffic to Router R4. Label 299856 is traffic arriving at Router 0 that corresponds to multicast LDP traffic from the root node R5 to the leaf egress nodes, R3 and R4.

As seen in the **show route detail** command output, the label operations for the protection path are the same for unicast LDP and multicast LDP traffic. The incoming LDP label at Router R0 is swapped to the LDP label advertised by Router R1 to Router R0. The RSVP label 299872 for the bypass tunnel is then pushed onto the packet. Penultimate hop-popping is used on the bypass tunnel, causing Router R2 to pop that label. Thus the packet arrives at Router R1 with the LDP label that it had originally advertised to Router R0.

Figure 56 on page 563 illustrates the label operation for unicast LDP and multicast LDP traffic protected by the RSVP bypass tunnel. The grey and blue boxes represent label values used under normal conditions for unicast and multicast LDP traffic, respectively. The dotted boxes represent label values used when the R0-R1 link fails.



Figure 56: LDP Link Protection Label Operation

**Case C: Dynamic RSVP LSP Only**

In this mode of operation, LFA is not used at all. A dynamic RSVP bypass LSP is automatically created in order to provide link protection. The output from the **show route detail** command and the label operations are similar to Case B, LFA and dynamic RSVP LSP mode.

**Sample Multicast LDP Link Protection Configuration**

To enable multicast LDP link protection, the following configuration is required on Router R0:



**NOTE:** In this sample, multicast LDP link protection is enabled on the ge-1/0/0 interface of Router R0 that connects to Router R1, although typically all the interfaces need to be configured for link protection.

```
Router R0 protocols {
 rsvp {
 interface all;
 interface ge-0/0/0.0 {
 disable;
 }
 }
 mpls {
 interface all;
 interface ge-0/0/0.0 {
 disable;
 }
 }
 ospf {
 traffic-engineering;
 }
}
```

```

 area 0.0.0.0 {
 interface lo0.0;
 interface ge-0/0/1.0 {
 link-protection;
 }
 interface ge-0/0/2.0;
 interface ge-0/0/3.0;
 }
}
ldp {
 make-before-break {
 timeout seconds;
 switchover-delay seconds;
 }
 interface ge-1/1/0.0 {
 link-protection {
 disable;
 dynamic-rsvp-lsp;
 }
 }
}
}

```

The following configuration statements apply to the different modes of multicast LDP protection as follows:

- **link-protection** statement at **[edit protocols ospf interface ge-0/0/1.0]**

This configuration is applied only for Case A (LFA only) and Case B (LFA and dynamic RSVP LSP) modes of multicast LDP link protection. Configuring link protection under an IGP is not required for Case C (dynamic RSVP LSP only).

- **link-protection** statement at **[edit protocols ldp interface ge-0/0/1.0]**

This configuration is required for all modes of multicast LDP protection. However, if the only LDP traffic present is unicast, and dynamic RSVP bypasses are not required, then this configuration is not required, as the **link-protection** statement at the **[edit protocols ospf interface ge-0/0/1.0]** hierarchy level results in LFA action for the LDP unicast traffic.

- **dynamic-rsvp-lsp** statement at **[edit protocols ldp interface ge-0/0/1.0 link-protection]**

This configuration is applied only for Case B (LFA and dynamic RSVP LSP) and Case C (dynamic RSVP LSP only) modes of LDP link protection. Dynamic RSVP LSP configuration does not apply to Case A (LFA only).

### Make-Before-Break

The make-before-break feature is enabled by default on Junos OS and provides some benefits for point-to-multipoint LSPs.

For a point-to-multipoint LSP, a label-switched router (LSR) selects the LSR that is its next hop to the root of the LSP as its upstream LSR. When the best path to reach the root changes, the LSR chooses a new upstream LSR. During this period, the LSP might be temporarily broken, resulting in packet loss until the LSP reconverges to a new

upstream LSR. The goal of make-before-break in this case is to minimize the packet loss. In cases where the best path from the LSR to the root changes but the LSP continues to forward traffic to the previous next hop to the root, a new LSP should be established before the old LSP is withdrawn to minimize the duration of packet loss.

Taking for example, after a link failure, a downstream LSR (for instance, LSR-D) still receives and/or forwards packets to the other downstream LSRs, as it continues to receive packets from the one hop RSVP LSP. Once routing converges, LSR-D selects a new upstream LSR (LSR-U) for this point-to-multipoint LSP's FEC (FEC-A). The new LSR might already be forwarding packets for FEC-A to the downstream LSRs other than LSR-D. After LSR-U receives a label for FEC-A from LSR-D, it notifies LSR-D when it has learnt that LSP for FEC-A has been established from the root to itself. When LSR-D receives such a notification, it changes its next hop for the LSP root to LSR-U. This is a route delete and add operation on LSR-D. At this point, LSR-D does an LSP switchover, and traffic tunneled through RSVP LSP or LFA is dropped, and traffic from LSR-U is accepted. The new transit route for LSR-U is added. The RPF check is changed to accept traffic from LSR-U and to drop traffic from the old upstream LSR, or the old route is deleted and the new route is added.

The assumption is that LSR-U has received a make-before-break notification from its upstream router for the FEC-A point-to-multipoint LSP and has installed a forwarding state for the LSP. At that point it should signal LSR-D by means of make-before-break notification that it has become part of the tree identified by FEC-A and that LSR-D should initiate its switchover to the LSP. Otherwise, LSR-U should remember that it needs to send notification to LSR-D when it receives a make-before-break notification from the upstream LSR for FEC-A and installs a forwarding state for this LSP. LSR-D continues to receive traffic from the old next hop to the root node using one hop RSVP LSP or LFA path until it switches over to the new point-to-multipoint LSP to LSR-U.

The make-before-break functionality with multicast LDP link protection includes the following features:

- Make-before-break capability

An LSR advertises that it is capable of handling make-before-break LSPs using the capability advertisement. If the peer is not make-before-break capable, the make-before-break parameters are not sent to this peer. If an LSR receives a make-before-break parameter from a downstream LSR (LSR-D) but the upstream LSR (LSR-U) is not make-before-break capable, the LSR immediately sends a make-before-break notification to LSR-D, and the make-before-break capable LSP is not established. Instead, the normal LSP is established.

- Make-before-break status code

The make-before-break status code includes:

- 1—make-before-break request
- 2—make-before-break acknowledgment

When a downstream LSR sends a label-mapping message for point-to-multipoint LSP, it includes the make-before-break status code as 1 (request). When the upstream LSR updates the forwarding state for the point-to-multipoint LSP, it informs the

downstream LSR with a notification message containing the make-before-break status code as 2 (acknowledgment). At that point, the downstream LSR does an LSP switchover.

### Caveats and Limitations

---

The Junos OS implementation of the LDP link protection feature has the following caveats and limitations:

- Make-before-break is not supported for the following point-to-multipoint LSPs on an egress LSR:
  - Next-generation multicast virtual private network (MVPN) with virtual routing and forwarding (VRF) label
  - Static LSP
- The following features are not supported:
  - Nonstop active routing for point-to-multipoint LSP in Junos OS Releases 12.3, 13.1 and 13.2
  - Graceful restart switchover point-to-multipoint LSP
  - Link protection for routing instance

## Understanding Multicast-Only Fast Reroute

---

Multicast-only fast reroute (MoFRR) minimizes packet loss in a network when there is a link failure. It works by enhancing multicast routing protocols like Protocol Independent Multicast (PIM) and multipoint Label Distribution Protocol (multipoint LDP). MoFRR is supported on MX Series routers with MPC line cards. As a prerequisite, the router must be set to **network-services enhanced-ip** mode, and all the line-cards in the router must be MPCs.

With MoFRR enabled, join messages are sent on primary and backup upstream paths. Data packets are received from both the primary path and the backup paths. The redundant packets are discarded based on priority (weights that are assigned to the primary and backup paths). When a failure is detected on the primary path, the repair is made by changing the interface on which packets are accepted to the secondary interface. Because the repair is local, it is fast—greatly improving convergence times in the event of a link failure on the primary path.

Currently, the most likely real-world application for MoFRR is streaming IPTV. IPTV streams are multicast as UDP streams. Therefore, any lost packets are not retransmitted, and this can result in a less-than-satisfactory user experience. MoFRR can be used to improve this situation.

When fast reroute is applied to unicast streams, an upstream router preestablishes MPLS label-switched paths (LSPs) or precomputes an IP loop-free alternate (LFA) fast reroute backup path to handle failure of a segment in the downstream path.

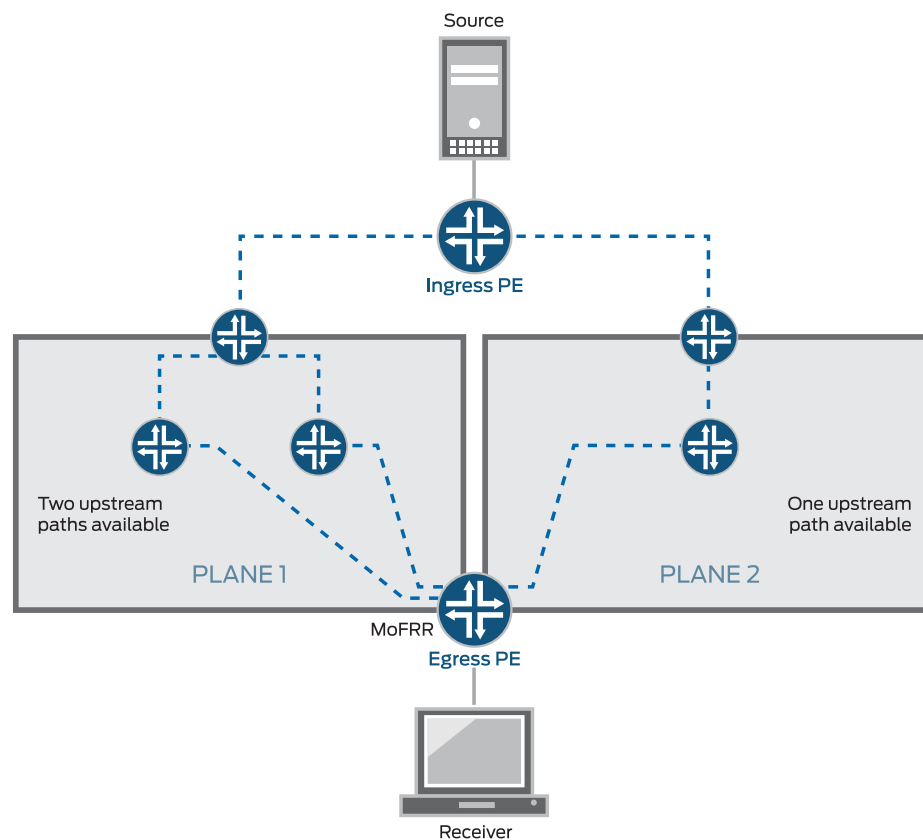
In multicast routing, the traffic distribution graphs are usually originated by the receiver. This is unlike unicast routing, which usually establishes the path from the source to the receiver. Protocols that are capable of establishing multicast distribution graphs are PIM (for IP), multipoint LDP (for MPLS), and RSVP-TE (for MPLS). Of these, PIM and multipoint LDP receivers initiate the distribution graph setup, and therefore these are the two multicast protocols for which MoFRR is supported.

In a multicast tree, performing a reactive repair upon detection of a network-component failure can lead to significant traffic loss due to delay in setting up the alternative path. MoFRR reduces traffic loss in a multicast distribution tree when a network component fails. With MoFRR, one of the downstream routers that supports this feature sets up an alternative path toward the source to receive a backup live stream of the same multicast traffic. When a failure is detected on the primary stream, the MoFRR router switches to the backup stream.

With MoFRR enabled, for each (S,G) entry, two of the available upstream interfaces are used to send a join message and to receive multicast traffic. The protocol attempts to select two disjoint paths if two such paths are available. If disjoint paths are not available, the protocol selects two nondisjoint paths. If two nondisjoint paths are not available, only a primary path is selected with no backup. MoFRR is supported for both IPv4 and IPv6 protocol families.

[Figure 57 on page 568](#) shows two paths from the egress provider edge (PE) router to the ingress PE router.

Figure 57: MoFRR Sample Topology



8041674

When enabled with MoFRR functionality, the egress router sets up two multicast trees, a primary path and a backup path, toward the multicast source for each (S,G). In other words, the egress router propagates the same (S,G) join messages toward two different upstream neighbors, thus creating two multicast trees.

One of the multicast trees goes through plane 1 and the other through plane 2, as shown in [Figure 57 on page 568](#). For each (S,G), the egress PE router forwards traffic received on the primary path and drops traffic received on the backup path.

MoFRR is supported on both equal-cost multipath (ECMP) paths and non-ECMP paths. Unicast loop-free alternate (LFA) routes need to be enabled to support MoFRR on non-ECMP paths. LFA routes are enabled with the **link-protection** statement in the interior gateway protocol (IGP) configuration. When you enable link protection on an OSPF or IS-IS interface, Junos OS creates a backup LFA path to the primary next hop for all destination routes that traverse the protected interface.

Junos OS implements MoFRR in the IP network for IP MoFRR and at the MPLS label-edge router (LER) for multipoint LDP MoFRR.

Multipoint LDP MoFRR is used at the egress node of an MPLS network, where the packets are forwarded to an IP network. In the case of multipoint LDP MoFRR, the two paths toward the upstream PE router are established for receiving two streams of MPLS packets at the LER. One of the streams (the primary) is accepted, and the other one (the backup)

is dropped at the LER. The backup stream is accepted if the primary path fails. A prerequisite for this feature is inband signaling support, as described in [“Understanding Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs” on page 598](#).

## PIM Functionality

Junos OS supports MoFRR for shortest-path tree (SPT) joins in PIM source-specific multicast (SSM) and any-source multicast (ASM). MoFRR is supported for both SSM and ASM ranges. To enable MoFRR for (\*G) joins, the `mofrr-asm-starg` configuration statement needs to be included. For each group G, either (S,G) or (\*G) (not both) will undergo MoFRR. (S,G) always takes precedence over (\*G).

With MoFRR enabled, a PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request. Preference is given to two paths that do not converge to the same immediate upstream router. PIM installs appropriate multicast routes with upstream RPF next hops with two (primary and backup) interfaces.

When the primary path fails, the backup path is upgraded to primary, and traffic is forwarded accordingly. If there are alternate paths available, a new backup path is calculated and the appropriate multicast route is updated or installed.

MoFRR can be enabled along with PIM join load balancing (with the `join-load-balance automatic` statement). However, in such cases the distribution of join messages among the links might not be even. When a new ECMP link is added, join messages on the primary path are redistributed and load-balanced. The join messages on the backup path might still follow the same path and might not be evenly redistributed.

MoFRR is enabled with a `[edit routing-options multicast stream-protection]` configuration and is managed by a set of filter policies. When an egress PIM router receives a join message or an IGMP report, the router checks for the MoFRR configuration.

If the MoFRR configuration is not present, PIM sends a join message upstream toward one upstream neighbor (for example, plane 2 in [Figure 57 on page 568](#)).

If the MoFRR configuration is present, Junos OS checks for a policy configuration.

If a policy is not present, Junos OS checks for primary and backup paths (upstream interfaces), and takes the following actions:

- If primary and backup paths are not available—PIM sends a join message upstream toward one upstream neighbor (for example, plane 2 in [Figure 57 on page 568](#)).
- If primary and backup paths are available—PIM sends the join message upstream toward two of the available upstream neighbors. Junos OS sets up primary and secondary multicast paths to receive multicast traffic (for example, plane 1 in [Figure 57 on page 568](#)).

If a policy is present, Junos OS checks whether the policy allows MoFRR for this (S,G), and takes the following actions:

- If the policy check fails—PIM sends a join message upstream toward one upstream neighbor (for example, plane 2 in [Figure 57 on page 568](#)).

- If the policy check passes—Junos OS checks for primary and backup paths (upstream interfaces).
- If the primary and backup paths are not available, PIM sends a join message upstream toward one upstream neighbor (for example, plane 2 in [Figure 57 on page 568](#)).
- If the primary and backup paths are available, PIM sends the join message upstream toward two of the available upstream neighbors. Junos OS sets up primary and secondary multicast paths to receive multicast traffic (for example, plane 1 in [Figure 57 on page 568](#)).

## Multipoint LDP Functionality

To avoid MPLS traffic duplication, the usual implementation of multipoint LDP selects only one upstream path. (See section 2.4.1.1. Determining One's 'upstream LSR' in RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*.)

For multipoint LDP MoFRR, the multipoint LDP node selects two separate upstream peers and sends two separate labels, one to each upstream peer. The same algorithm described in RFC 6388 is used to select the primary upstream path. The backup upstream path selection again uses the same algorithm but excludes the primary upstream LSR as a candidate. Two streams of MPLS traffic are sent to the egress node from the two different upstream peers. The MPLS traffic from only one of the upstream neighbors is selected as the primary path to accept the traffic, and the other path becomes the backup path. The traffic on the backup path is dropped. When the primary upstream path fails, the traffic from the backup path is then accepted. The multipoint LDP node selects the two upstream paths based on the interior gateway protocol (IGP) root node next hop.

A forwarding equivalency class (FEC) is a group of IP packets that are forwarded in the same manner, over the same path, and with the same forwarding treatment. Normally, the label that is put on a particular packet represents the FEC to which that packet is assigned. In MoFRR, two routes are placed into the mpls.0 table for each FEC—one route for the primary label and the other route for the backup label.

If there are parallel links toward the same immediate upstream node, both parallel links are considered to be the primary. At any point in time, the upstream node sends traffic on only one of the multiple parallel links.

A bud node is an LSR that is an egress LSR, but also has one or more directly connected downstream LSRs. In the case of a bud node, the traffic from the primary upstream path is forwarded to a downstream LSR. If the primary upstream path fails, the MPLS traffic from the backup upstream path is forwarded to the downstream LSR. This means that the downstream LSR next hop is added to both MPLS routes along with the egress next hop.

MoFRR for multipoint LDP is enabled with a **[edit routing-options multicast stream-protection]** configuration and is managed by a set of filter policies.

If the multipoint LDP point-to-multipoint FEC is enabled for MoFRR, the following additional considerations are factored into upstream path selection:



- The targeted LDP sessions are skipped if there is a nontargeted LDP session. If there is a single targeted LDP session, the targeted LDP session is selected, but the corresponding point-to-multipoint FEC loses the MoFRR capability because there is no interface associated with the targeted LDP session.
- All interfaces that belong to the same upstream LSR are considered to be the primary path.
- For any root-node route updates, the upstream path is changed based on the latest next hops from the IGP. If a better path is available, multipoint LDP attempts to switch to the better path.


## Packet Forwarding

For both PIM and multipoint LDP, multicast source stream selection is performed at the incoming interface. This prevents duplicate streams from being sent across the fabric and prevents multiple route lookups that result in drops, thus preserving fabric bandwidth and maximizing forwarding performance.

For PIM, each IP multicast stream contains the same destination address. Regardless of the interface on which the packets arrive, the packets have the same route. An interface list is attached to the route. Junos OS checks the interface upon which each packet arrives and forwards only those that are from the primary interface. If the interface matches a secondary interface, the packets are dropped. If no match is found, the packets are handled as exceptions in the control plane.

This process is shown in [Figure 58 on page 571](#).

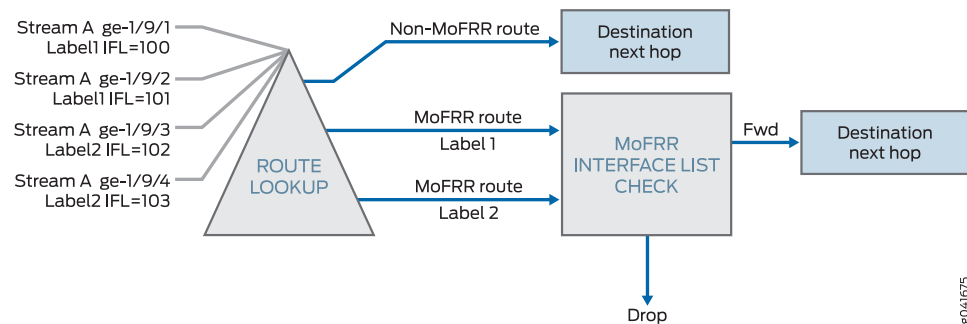
**Figure 58: MoFRR IP Route Lookup in the Packet Forwarding Engine**



For multipoint LDP, multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.

This process is shown in [Figure 59 on page 572](#).

Figure 59: MoFRR MPLS Route Lookup in the Packet Forwarding Engine



## Limitations and Caveats

MoFRR has the following limitations and caveats:

- MoFRR failure detection is supported for immediate link protection of the router on which MoFRR is enabled and not on all the links (end-to-end) in the multicast traffic path.
- MoFRR supports FRR on two selected disjoint paths toward the source. Two of the selected upstream neighbors cannot be on the same interface—in other words, two upstream neighbors on a LAN segment. The same is true if the upstream interface happens to be a multicast tunnel interface.
- Detection of the maximum end-to-end disjoint upstream paths is not supported. The egress router only makes sure that there is a disjoint upstream node (the immediate previous hop). PIM and multipoint LDP do not support the equivalent of explicit route objects (EROs). Hence, disjoint upstream path detection is limited to control over the immediately previous hop node. Because of this limitation, the path to the upstream node of the previous hop selected as primary and backup might be shared.
- MoFRR does not apply to multipoint LDP traffic received on an RSVP tunnel because the RSVP tunnel is not associated with any interface.
- Some traffic loss is seen in the following scenarios:
  - A better upstream path becomes available on an egress node.
  - MoFRR is enabled or disabled on the egress node while there is an active traffic stream flowing.
- PIM join load balancing for join messages for backup paths are not supported.
- For a multicast group G, MoFRR is not allowed for both (S,G) and (\*,G) join messages. (S,G) join messages have precedence over (\*,G).
- MoFRR is not supported for multicast traffic streams that use two different multicast groups. Each (S,G) combination is treated as a unique multicast traffic stream.
- The bidirectional PIM range is not supported for MoFRR.
- PIM dense-mode is not supported for MoFRR.

- Mixed upstream MoFRR is not supported. This refers to PIM multipoint LDP in-band signaling, wherein one upstream path is through multipoint LDP and the second upstream path is through PIM.
- Multicast statistics for the backup traffic stream are not maintained by PIM and therefore are not available in the operational output of **show** commands.
- Multipoint LDP labels as inner labels are not supported.
- If the source is reachable through multiple ingress provider edge (PE) routers, multipoint LDP MoFRR is not supported.
- Targeted upstream sessions are not selected as the upstream node for MoFRR.
- Rate monitoring is not supported.
- Multipoint LDP link protection on the backup path is not supported because there is no support for MoFRR inner labels.

**Related  
Documentation**

- [Configuring Multicast-Only Fast Reroute on page 573](#)
- *Example: Configuring Multicast-Only Fast Reroute in a PIM Domain*
- [Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain on page 576](#)

## Configuring Multicast-Only Fast Reroute

You can configure multicast-only fast reroute (MoFRR) to minimize packet loss in a network when there is a link failure.

When fast reroute is applied to unicast streams, an upstream router preestablishes MPLS label-switched paths (LSPs) or precomputes an IP loop-free alternate (LFA) fast reroute backup path to handle failure of a segment in the downstream path.

In multicast routing, the traffic distribution graphs are usually originated by the receiver. This is unlike unicast routing, which usually establishes the path from the source to the receiver. Protocols that are capable of establishing multicast distribution graphs are PIM (for IP), multipoint LDP (for MPLS) and RSVP-TE (for MPLS). Of these, PIM and multipoint LDP receivers initiate the distribution graph setup, and therefore these are the two multicast protocols for which MoFRR is supported.

MoFRR is supported on MX Series routers with MPC line cards. As a prerequisite, all the line cards in the router must be MPCs.

Make sure that all of the

To configure MoFRR:

1. Set the router to enhanced IP mode.

```
[edit chassis]
user@host# set network-services enhanced-ip
```

2. Enable MoFRR.

```
[edit routing-options multicast]
user@host# set stream-protection
```

3. (Optional) Configure a routing policy that filters for a restricted set of multicast streams to be affected by your MoFRR configuration.

You can apply filters that are based on source or group addresses.

For example:

```
[edit policy-options]
policy-statement mofrr-select {
 term A {
 from {
 source-address-filter 225.1.1.1/32 exact;
 }
 then {
 accept;
 }
 }
 term B {
 from {
 source-address-filter 226.0.0.0/8 orlonger;
 }
 then {
 accept;
 }
 }
 term C {
 from {
 source-address-filter 227.1.1.0/24 orlonger;
 source-address-filter 227.4.1.0/24 orlonger;
 source-address-filter 227.16.1.0/24 orlonger;
 }
 then {
 accept;
 }
 }
 term D {
 from {
 source-address-filter 227.1.1.1/32 exact
 }
 then {
 reject; #MoFRR disabled
 }
 }
 ...
}
```

4. (Optional) If you configured a routing policy to filter the set of to be affected by your MoFRR configuration, apply the policy.

```
[edit routing-options multicast stream-protection]
user@host# set policy policy-name
```

For example:

```
routing-options {
```

```

multicast {
 stream-protection {
 policy mofrr-select
 }
}

```

5. (Optional) In a PIM domain with MoFRR, allow MoFRR to be applied to any-source multicast (ASM) (\*G) joins.

This is not supported for multipoint LDP MoFRR.

```

[edit routing-options multicast stream-protection]
user@host# set mofrr-asm-starg

```

6. (Optional) In a PIM domain with MoFRR, allow only a disjoint RPF (an RPF on a separate plane) to be selected as the backup RPF path.

This is not supported for multipoint LDP MoFRR. In a multipoint LDP MoFRR domain, the same label is shared between parallel links to the same upstream neighbor. This is not the case in a PIM domain, where each link forms a neighbor. The

**mofrr-disjoint-upstream-only** statement does not allow a backup RPF path to be selected if the path goes to the same upstream neighbor as that of the primary RPF path. This ensures that MoFRR is triggered only on a topology that has multiple RPF upstream neighbors.

```

[edit routing-options multicast stream-protection]
user@host# set mofrr-disjoint-upstream-only

```

7. (Optional) In a PIM domain with MoFRR, prevent sending join messages on the backup path, but retain all other MoFRR functionality.

This is not supported for multipoint LDP MoFRR.

```

[edit routing-options multicast stream-protection]
user@host# set mofrr-no-backup-join

```

8. (Optional) In a PIM domain with MoFRR, allow new primary path selection to be based on the unicast gateway selection for the unicast route to the source and to change when there is a change in the unicast selection, rather than having the backup path be promoted as primary. This ensures that the primary RPF hop is always on the best path.

When you include the **mofrr-primary-selection-by-routing** statement, the backup path is not guaranteed to get promoted to be the new primary path when the primary path goes down.

This is not supported for multipoint LDP MoFRR.

```

[edit routing-options multicast stream-protection]
user@host# set mofrr-primary-selection-by-routing

```

#### Related Documentation

- [Understanding Multicast-Only Fast Reroute on page 566](#)
- [Example: Configuring Multicast-Only Fast Reroute in a PIM Domain](#)
- [Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain on page 576](#)

## Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain

---

This example shows how to configure multicast-only fast reroute (MoFRR) to minimize packet loss in a network when there is a link failure.

Multipoint LDP MoFRR is used at the egress node of an MPLS network, where the packets are forwarded to an IP network. In the case of multipoint LDP MoFRR, the two paths toward the upstream provider edge (PE) router are established for receiving two streams of MPLS packets at the label-edge router (LER). One of the streams (the primary) is accepted, and the other one (the backup) is dropped at the LER. The backup stream is accepted if the primary path fails.

- [Requirements on page 576](#)
- [Overview on page 576](#)
- [CLI Quick Configuration on page 577](#)
- [Configuration on page 583](#)
- [Verification on page 588](#)

### Requirements

No special configuration beyond device initialization is required before configuring this example.

In a multipoint LDP domain, for MoFRR to work, only the egress PE router needs to have MoFRR enabled. The other routers do not need to support MoFRR.

MoFRR is supported on MX Series platforms with MPC line cards. As a prerequisite, the router must be set to **network-services enhanced-ip** mode, and all the line-cards in the platform must be MPCs.

This example requires Junos OS Release 14.1 or later on the egress PE router.

### Overview

In this example, Device R3 is the egress edge router. MoFRR is enabled on this device only.

OSPF is used for connectivity, though any interior gateway protocol (IGP) or static routes can be used.

For testing purposes, routers are used to simulate the source and the receiver. Device R4 and Device R8 are configured to statically join the desired group by using the **set protocols igmp interface interface-name static group group** command. In the case when a real multicast receiver host is not available, as in this example, this static IGMP configuration is useful. On the receivers, to make them listen to the multicast group address, this example uses **set protocols sap listen group**.

MoFRR configuration includes a policy option that is not shown in this example, but is explained separately. The option is configured as follows:

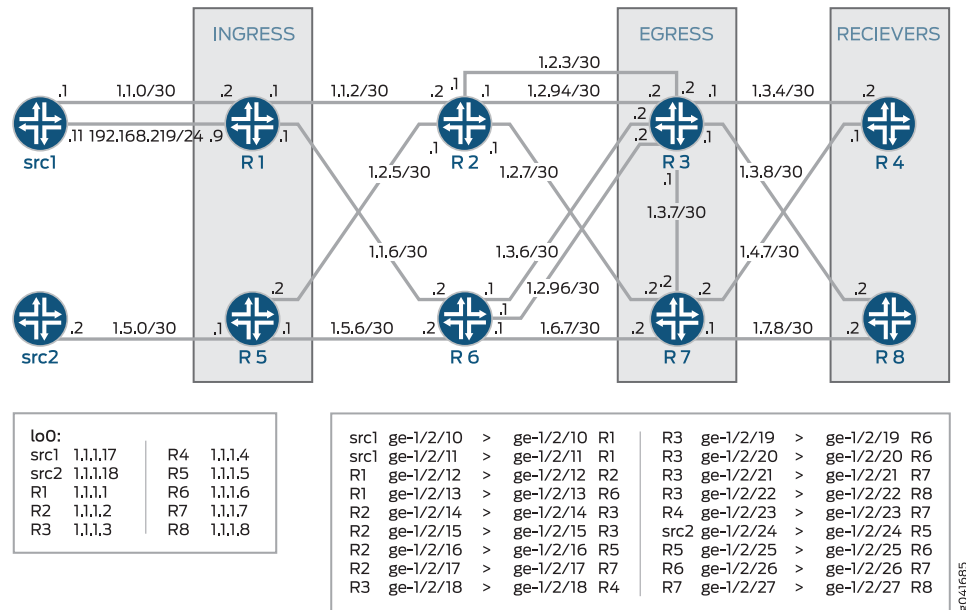
```
stream-protection {
```

```
policy policy-name;
```

## Topology

Figure 60 on page 577 shows the sample network.

### Figure 60: MoFRR in a Multipoint LDP Domain



“CLI Quick Configuration” on page 577 shows the configuration for all of the devices in Figure 60 on page 577.

The section “Configuration” on page 583 describes the steps on Device R3.

## CLI Quick Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
Device src1
set interfaces ge-1/2/10 unit 0 description src1-to-R1
set interfaces ge-1/2/10 unit 0 family inet address 1.1.0.1/30
set interfaces ge-1/2/11 unit 0 description src1-to-R1
set interfaces ge-1/2/11 unit 0 family inet address 192.168.219.11/24
set interfaces lo0 unit 0 family inet address 1.1.1.17/32
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

```
Device src2 set interfaces ge-1/2/24 unit 0 description src2-to-R5
 set interfaces ge-1/2/24 unit 0 family inet address 1.5.0.2/30
 set interfaces lo0 unit 0 family inet address 1.1.1.18/32
 set protocols rsvp interface all
 set protocols ospf area 0.0.0.0 interface all
 set protocols ospf area 0.0.0.0 interface lo0.0 passive
```

**Device R1**

```
set interfaces ge-1/2/12 unit 0 description R1-to-R2
set interfaces ge-1/2/12 unit 0 family inet address 1.1.2.1/30
set interfaces ge-1/2/12 unit 0 family mpls
set interfaces ge-1/2/13 unit 0 description R1-to-R6
set interfaces ge-1/2/13 unit 0 family inet address 1.1.6.1/30
set interfaces ge-1/2/13 unit 0 family mpls
set interfaces ge-1/2/10 unit 0 description R1-to-src1
set interfaces ge-1/2/10 unit 0 family inet address 1.1.0.2/30
set interfaces ge-1/2/11 unit 0 description R1-to-src1
set interfaces ge-1/2/11 unit 0 family inet address 192.168.219.9/30
set interfaces lo0 unit 0 family inet address 1.1.1.1/32
set protocols rsvp interface all
set protocols mpls interface all
set protocols bgp group ibgp local-address 1.1.1.1
set protocols bgp group ibgp export static-route-tobgp
set protocols bgp group ibgp peer-as 10
set protocols bgp group ibgp neighbor 1.1.1.3
set protocols bgp group ibgp neighbor 1.1.1.7
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-1/2/12.0
set protocols ldp interface ge-1/2/13.0
set protocols ldp interface lo0.0
set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim rp static address 1.1.1.5
set protocols pim interface lo0.0
set protocols pim interface ge-1/2/10.0
set protocols pim interface ge-1/2/11.0
set policy-options policy-statement mldppim-ex term B from source-address-filter
192.168.0.0/24 orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address
1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term A from source-address-filter
1.1.1.7/32 orlonger
set policy-options policy-statement mldppim-ex term A from source-address-filter
1.1.0.0/30 orlonger
set policy-options policy-statement mldppim-ex term A then accept
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set routing-options autonomous-system 10
```

**Device R2**

```
set interfaces ge-1/2/12 unit 0 description R2-to-R1
set interfaces ge-1/2/12 unit 0 family inet address 1.1.2.2/30
set interfaces ge-1/2/12 unit 0 family mpls
set interfaces ge-1/2/14 unit 0 description R2-to-R3
set interfaces ge-1/2/14 unit 0 family inet address 1.2.3.1/30
set interfaces ge-1/2/14 unit 0 family mpls
set interfaces ge-1/2/16 unit 0 description R2-to-R5
set interfaces ge-1/2/16 unit 0 family inet address 1.2.5.1/30
```



```

set interfaces ge-1/2/16 unit 0 family mpls
set interfaces ge-1/2/17 unit 0 description R2-to-R7
set interfaces ge-1/2/17 unit 0 family inet address 1.2.7.1/30
set interfaces ge-1/2/17 unit 0 family mpls
set interfaces ge-1/2/15 unit 0 description R2-to-R3
set interfaces ge-1/2/15 unit 0 family inet address 1.2.94.1/30
set interfaces ge-1/2/15 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.2/32
set interfaces lo0 unit 0 family mpls
set protocols rsvp interface all
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface all
set protocols ldp p2mp
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.0.0/24 orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address
 1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set routing-options autonomous-system 10

```

**Device R3**

```

set chassis network-services enhanced-ip
set interfaces ge-1/2/14 unit 0 description R3-to-R2
set interfaces ge-1/2/14 unit 0 family inet address 1.2.3.2/30
set interfaces ge-1/2/14 unit 0 family mpls
set interfaces ge-1/2/18 unit 0 description R3-to-R4
set interfaces ge-1/2/18 unit 0 family inet address 1.3.4.1/30
set interfaces ge-1/2/18 unit 0 family mpls
set interfaces ge-1/2/19 unit 0 description R3-to-R6
set interfaces ge-1/2/19 unit 0 family inet address 1.3.6.2/30
set interfaces ge-1/2/19 unit 0 family mpls
set interfaces ge-1/2/21 unit 0 description R3-to-R7
set interfaces ge-1/2/21 unit 0 family inet address 1.3.7.1/30
set interfaces ge-1/2/21 unit 0 family mpls
set interfaces ge-1/2/22 unit 0 description R3-to-R8
set interfaces ge-1/2/22 unit 0 family inet address 1.3.8.1/30
set interfaces ge-1/2/22 unit 0 family mpls
set interfaces ge-1/2/15 unit 0 description R3-to-R2
set interfaces ge-1/2/15 unit 0 family inet address 1.2.94.2/30
set interfaces ge-1/2/15 unit 0 family mpls
set interfaces ge-1/2/20 unit 0 description R3-to-R6
set interfaces ge-1/2/20 unit 0 family inet address 1.2.96.2/30
set interfaces ge-1/2/20 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.3/32 primary
set routing-options autonomous-system 10
set routing-options multicast stream-protection
set protocols rsvp interface all
set protocols mpls interface all
set protocols bgp group ibgp local-address 1.1.1.3
set protocols bgp group ibgp peer-as 10
set protocols bgp group ibgp neighbor 1.1.1.1
set protocols bgp group ibgp neighbor 1.1.1.5

```

```

set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface all
set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface lo0.0
set protocols pim interface ge-1/2/18.0
set protocols pim interface ge-1/2/22.0
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.0.0/24 orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term A from source-address-filter
 1.1.0.1/30 orlonger
set policy-options policy-statement mldppim-ex term A then accept
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept

```

**Device R4**

```

set interfaces ge-1/2/18 unit 0 description R4-to-R3
set interfaces ge-1/2/18 unit 0 family inet address 1.3.4.2/30
set interfaces ge-1/2/18 unit 0 family mpls
set interfaces ge-1/2/23 unit 0 description R4-to-R7
set interfaces ge-1/2/23 unit 0 family inet address 1.4.7.1/30
set interfaces lo0 unit 0 family inet address 1.1.1.4/32
set protocols igmp interface ge-1/2/18.0 version 3
set protocols igmp interface ge-1/2/18.0 static group 232.1.1.1 group-count 2
set protocols igmp interface ge-1/2/18.0 static group 232.1.1.1 source 192.168.219.11
set protocols igmp interface ge-1/2/18.0 static group 232.2.2.2 source 1.2.7.7
set protocols sap listen 232.1.1.1
set protocols sap listen 232.2.2.2
set protocols rsvp interface all
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface ge-1/2/23.0
set protocols pim interface ge-1/2/18.0
set protocols pim interface lo0.0
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.0.0/24 orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address
 1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set routing-options autonomous-system 10

```

**Device R5**

```

set interfaces ge-1/2/24 unit 0 description R5-to-src2
set interfaces ge-1/2/24 unit 0 family inet address 1.5.0.1/30
set interfaces ge-1/2/16 unit 0 description R5-to-R2
set interfaces ge-1/2/16 unit 0 family inet address 1.2.5.2/30
set interfaces ge-1/2/16 unit 0 family mpls
set interfaces ge-1/2/25 unit 0 description R5-to-R6
set interfaces ge-1/2/25 unit 0 family inet address 1.5.6.1/30
set interfaces ge-1/2/25 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.5/32
set protocols rsvp interface all
set protocols mpls interface all
set protocols bgp group ibgp local-address 1.1.1.5
set protocols bgp group ibgp export static-route-tobgp
set protocols bgp group ibgp peer-as 10
set protocols bgp group ibgp neighbor 1.1.1.7
set protocols bgp group ibgp neighbor 1.1.1.3
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-1/2/16.0
set protocols ldp interface ge-1/2/25.0
set protocols ldp p2mp
set protocols pim interface lo0.0
set protocols pim interface ge-1/2/24.0
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set routing-options autonomous-system 10

```

**Device R6**

```

set interfaces ge-1/2/13 unit 0 description R6-to-R1
set interfaces ge-1/2/13 unit 0 family inet address 1.1.6.2/30
set interfaces ge-1/2/13 unit 0 family mpls
set interfaces ge-1/2/19 unit 0 description R6-to-R3
set interfaces ge-1/2/19 unit 0 family inet address 1.3.6.1/30
set interfaces ge-1/2/19 unit 0 family mpls
set interfaces ge-1/2/25 unit 0 description R6-to-R5
set interfaces ge-1/2/25 unit 0 family inet address 1.5.6.2/30
set interfaces ge-1/2/25 unit 0 family mpls
set interfaces ge-1/2/26 unit 0 description R6-to-R7
set interfaces ge-1/2/26 unit 0 family inet address 1.6.7.1/30
set interfaces ge-1/2/26 unit 0 family mpls
set interfaces ge-1/2/20 unit 0 description R6-to-R3
set interfaces ge-1/2/20 unit 0 family inet address 1.2.96.1/30
set interfaces ge-1/2/20 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.6/30
set protocols rsvp interface all
set protocols mpls interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface all
set protocols ldp p2mp

```

**Device R7**

```

set interfaces ge-1/2/17 unit 0 description R7-to-R2
set interfaces ge-1/2/17 unit 0 family inet address 1.2.7.2/30

```

```

set interfaces ge-1/2/17 unit 0 family mpls
set interfaces ge-1/2/21 unit 0 description R7-to-R3
set interfaces ge-1/2/21 unit 0 family inet address 1.3.7.2/30
set interfaces ge-1/2/21 unit 0 family mpls
set interfaces ge-1/2/23 unit 0 description R7-to-R4
set interfaces ge-1/2/23 unit 0 family inet address 1.4.7.2/30
set interfaces ge-1/2/23 unit 0 family mpls
set interfaces ge-1/2/26 unit 0 description R7-to-R6
set interfaces ge-1/2/26 unit 0 family inet address 1.6.7.2/30
set interfaces ge-1/2/26 unit 0 family mpls
set interfaces ge-1/2/27 unit 0 description R7-to-R8
set interfaces ge-1/2/27 unit 0 family inet address 1.7.8.1/30
set interfaces ge-1/2/27 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.7/32
set protocols rsvp interface all
set protocols mpls interface all
set protocols bgp group ibgp local-address 1.1.1.7
set protocols bgp group ibgp export static-route-tobgp
set protocols bgp group ibgp peer-as 10
set protocols bgp group ibgp neighbor 1.1.1.5
set protocols bgp group ibgp neighbor 1.1.1.1
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols ldp interface ge-1/2/17.0
set protocols ldp interface ge-1/2/21.0
set protocols ldp interface ge-1/2/26.0
set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface lo0.0
set protocols pim interface ge-1/2/27.0
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.0.0/24 orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term A from source-address-filter
 1.1.0.1/30 orlonger
set policy-options policy-statement mldppim-ex term A then accept
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set routing-options autonomous-system 10
set routing-options multicast stream-protection policy mldppim-ex

```

**Device R8**

```

set interfaces ge-1/2/22 unit 0 description R8-to-R3
set interfaces ge-1/2/22 unit 0 family inet address 1.3.8.2/30
set interfaces ge-1/2/22 unit 0 family mpls
set interfaces ge-1/2/27 unit 0 description R8-to-R7
set interfaces ge-1/2/27 unit 0 family inet address 1.7.8.2/30
set interfaces ge-1/2/27 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.8/32
set protocols igmp interface ge-1/2/22.0 version 3
set protocols igmp interface ge-1/2/22.0 static group 232.1.1.1 group-count 2
set protocols igmp interface ge-1/2/22.0 static group 232.1.1.1 source 192.168.219.11
set protocols igmp interface ge-1/2/22.0 static group 232.2.2.2 source 1.2.7.7

```

```

set protocols sap listen 232.1.1.1
set protocols sap listen 232.2.2.2
set protocols rsvp interface all
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface lo0.0 passive
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface ge-1/2/27.0
set protocols pim interface ge-1/2/22.0
set protocols pim interface lo0.0
set policy-options policy-statement static-route-tobgp term static from protocol static
set policy-options policy-statement static-route-tobgp term static from protocol direct
set policy-options policy-statement static-route-tobgp term static then accept
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.0.0/24 orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address
 1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set routing-options autonomous-system 10

```

## Configuration

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R3:

1. Enable enhanced IP mode.
 

```

[edit chassis]
user@R3# set network-services enhanced-ip

```
2. Configure the device interfaces.
 

```

[edit interfaces]
user@R3# set ge-1/2/14 unit 0 description R3-to-R2
user@R3# set ge-1/2/14 unit 0 family inet address 1.2.3.2/30
user@R3# set ge-1/2/14 unit 0 family mpls

user@R3# set ge-1/2/18 unit 0 description R3-to-R4
user@R3# set ge-1/2/18 unit 0 family inet address 1.3.4.1/30
user@R3# set ge-1/2/18 unit 0 family mpls

user@R3# set ge-1/2/19 unit 0 description R3-to-R6
user@R3# set ge-1/2/19 unit 0 family inet address 1.3.6.2/30
user@R3# set ge-1/2/19 unit 0 family mpls

user@R3# set ge-1/2/21 unit 0 description R3-to-R7
user@R3# set ge-1/2/21 unit 0 family inet address 1.3.7.1/30
user@R3# set ge-1/2/21 unit 0 family mpls

user@R3# set ge-1/2/22 unit 0 description R3-to-R8

```

```
user@R3# set ge-1/2/22 unit 0 family inet address 1.3.8.1/30
user@R3# set ge-1/2/22 unit 0 family mpls
```

```
user@R3# set ge-1/2/15 unit 0 description R3-to-R2
user@R3# set ge-1/2/15 unit 0 family inet address 1.2.94.2/30
user@R3# set ge-1/2/15 unit 0 family mpls
```

```
user@R3# set ge-1/2/20 unit 0 description R3-to-R6
user@R3# set ge-1/2/20 unit 0 family inet address 1.2.96.2/30
user@R3# set ge-1/2/20 unit 0 family mpls
```

```
user@R3# set lo0 unit 0 family inet address 1.1.1.3/32 primary
```

3. Configure the autonomous system (AS) number.

```
user@R3# set routing-options autonomous-system 10
```

4. Configure the routing policies.

```
[edit policy-options policy-statement mldppim-ex]
user@R3# set term B from source-address-filter 192.168.0.0/24 orlonger
user@R3# set term B from source-address-filter 192.168.219.11/32 orlonger
user@R3# set term B then accept
user@R3# set term A from source-address-filter 1.1.0.1/30 orlonger
user@R3# set term A then accept
```

```
[edit policy-options policy-statement static-route-tobgp]
user@R3# set term static from protocol static
user@R3# set term static from protocol direct
user@R3# set term static then accept
```

5. Configure PIM.

```
[edit protocols pim]
user@R3# set mldp-inband-signalling policy mldppim-ex
user@R3# set interface lo0.0
user@R3# set interface ge-1/2/18.0
user@R3# set interface ge-1/2/22.0
```

6. Configure LDP.

```
[edit protocols ldp]
user@R3# set interface all
user@R3# set p2mp
```

7. Configure an IGP or static routes.

```
[edit protocols ospf]
user@R3# set traffic-engineering
user@R3# set area 0.0.0.0 interface all
user@R3# set area 0.0.0.0 interface fxp0.0 disable
user@R3# set area 0.0.0.0 interface lo0.0 passive
```

8. Configure internal BGP.

```
[edit protocols bgp group ibgp]
user@R3# set local-address 1.1.1.3
user@R3# set peer-as 10
user@R3# set neighbor 1.1.1.1
```

```
user@R3# set neighbor 1.1.1.5
```

9. Configure MPLS and, optionally, RSVP.

```
[edit protocols mpls]
user@R3# set interface all
```

```
[edit protocols rsvp]
user@R3# set interface all
```

10. Enable MoFRR.

```
[edit routing-options multicast]
user@R3# set stream-protection
```

**Results** From configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R3# show chassis
network-services enhanced-ip;
```

```
user@R3# show interfaces
```

```
ge-1/2/14 {
 unit 0 {
 description R3-to-R2;
 family inet {
 address 1.2.3.2/30;
 }
 family mpls;
 }
}
```

```
ge-1/2/18 {
 unit 0 {
 description R3-to-R4;
 family inet {
 address 1.3.4.1/30;
 }
 family mpls;
 }
}
```

```
ge-1/2/19 {
 unit 0 {
 description R3-to-R6;
 family inet {
 address 1.3.6.2/30;
 }
 family mpls;
 }
}
```

```
ge-1/2/21 {
 unit 0 {
 description R3-to-R7;
 family inet {
 address 1.3.7.1/30;
 }
 }
}
```

```
 }
 family mpls;
 }
}
ge-1/2/22 {
 unit 0 {
 description R3-to-R8;
 family inet {
 address 1.3.8.1/30;
 }
 family mpls;
 }
}
ge-1/2/15 {
 unit 0 {
 description R3-to-R2;
 family inet {
 address 1.2.94.2/30;
 }
 family mpls;
 }
}
ge-1/2/20 {
 unit 0 {
 description R3-to-R6;
 family inet {
 address 1.2.96.2/30;
 }
 family mpls;
 }
}
lo0 {
 unit 0 {
 family inet {
 address 192.168.15.1/32;
 address 1.1.1.3/32 {
 primary;
 }
 }
 }
}

user@R3# show protocols
rsvp {
 interface all;
}
mpls {
 interface all;
}
bgp {
 group ibgp {
 local-address 1.1.1.3;
 peer-as 10;
 neighbor 1.1.1.1;
 neighbor 1.1.1.5;
 }
}
```



```

}
ospf {
 traffic-engineering;
 area 0.0.0.0 {
 interface all;
 interface fxp0.0 {
 disable;
 }
 interface lo0.0 {
 passive;
 }
 }
}
ldp {
 interface all;
 p2mp;
}
pim {
 mldp-inband-signalling {
 policy mldppim-ex;
 }
 interface lo0.0;
 interface ge-1/2/18.0;
 interface ge-1/2/22.0;
}

user@R3# show policy-options
policy-statement mldppim-ex {
 term B {
 from {
 source-address-filter 192.168.0.0/24 orlonger;
 source-address-filter 192.168.219.11/32 orlonger;
 }
 then accept;
 }
 term A {
 from {
 source-address-filter 1.1.0.1/30 orlonger;
 }
 then accept;
 }
}
policy-statement static-route-tobgp {
 term static {
 from protocol [static direct];
 then accept;
 }
}

user@R3# show routing-options
autonomous-system 10;
multicast {
 stream-protection;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Checking the LDP Point-to-Multipoint Forwarding Equivalency Classes on page 588](#)
- [Examining the Label Information on page 588](#)
- [Checking the Multicast Routes on page 590](#)
- [Checking the LDP Point-to-Multipoint Traffic Statistics on page 591](#)

---

### Checking the LDP Point-to-Multipoint Forwarding Equivalency Classes

**Purpose** Make sure the MoFRR is enabled, and determine what labels are being used.

**Action** user@R3> show ldp p2mp fec

```
LDP P2MP FECs:
P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src: 192.168.219.11
 MoFRR enabled
 Fec type: Egress (Active)
 Label: 301568
P2MP root-addr 1.1.1.1, grp: 232.1.1.2, src: 192.168.219.11
 MoFRR enabled
 Fec type: Egress (Active)
 Label: 301600
```

**Meaning** The output shows that MoFRR is enabled, and it shows that the labels 301568 and 301600 are being used for the two multipoint LDP point-to-multipoint LSPs.

---

### Examining the Label Information

**Purpose** Make sure that the egress device has two upstream interfaces for the multicast group join.

**Action** user@R3> show route label 301568 detail

```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
 *LDP Preference: 9
 Next hop type: Flood
 Address: 0x2735208
 Next-hop reference count: 3
 Next hop type: Router, Next hop index: 1397
 Address: 0x2735d2c
 Next-hop reference count: 3
 Next hop: 1.3.8.2 via ge-1/2/22.0
 Label operation: Pop
 Load balance label: None;
 Next hop type: Router, Next hop index: 1395
 Address: 0x2736290
 Next-hop reference count: 3
 Next hop: 1.3.4.2 via ge-1/2/18.0
 Label operation: Pop
 Load balance label: None;
 State: <Active Int AckRequest MulticastRPF>
 Local AS: 10
 Age: 54:05 Metric: 1
 Validation State: unverified
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
 Primary Upstream : 1.1.1.3:0--1.1.1.2:0
 RPF Nexthops :
 ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
 ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
 Backup Upstream : 1.1.1.3:0--1.1.1.6:0
 RPF Nexthops :
 ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffff
 ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffff

```

user@R3> show route label 301600 detail

```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301600 (1 entry, 1 announced)
 *LDP Preference: 9
 Next hop type: Flood
 Address: 0x27356b4
 Next-hop reference count: 3
 Next hop type: Router, Next hop index: 1520
 Address: 0x27350f4
 Next-hop reference count: 3
 Next hop: 1.3.8.2 via ge-1/2/22.0
 Label operation: Pop
 Load balance label: None;
 Next hop type: Router, Next hop index: 1481
 Address: 0x273645c
 Next-hop reference count: 3
 Next hop: 1.3.4.2 via ge-1/2/18.0
 Label operation: Pop
 Load balance label: None;
 State: <Active Int AckRequest MulticastRPF>

```

```
Local AS: 10
Age: 54:25 Metric: 1
Validation State: unverified
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.2, src:
192.168.219.11
Primary Upstream : 1.1.1.3:0--1.1.1.6:0
RPF Nexthops :
 ge-1/2/20.0, 1.2.96.1, Label: 301600, weight: 0x1
 ge-1/2/19.0, 1.3.6.1, Label: 301600, weight: 0x1
Backup Upstream : 1.1.1.3:0--1.1.1.2:0
RPF Nexthops :
 ge-1/2/15.0, 1.2.94.1, Label: 301616, weight: 0xffffe
 ge-1/2/14.0, 1.2.3.1, Label: 301616, weight: 0xffffe
```

**Meaning** The output shows the primary upstream paths and the backup upstream paths. It also shows the RPF next hops.

---

### Checking the Multicast Routes

**Purpose** Examine the IP multicast forwarding table to make sure that there is an upstream RPF interface list, with a primary and a backup interface.

```

Action user@R3> show ldp p2mp path
P2MP path type: Transit/Egress
 Output Session (label): 1.1.1.2:0 (301568) (Primary)
 Egress Nexthops: Interface ge-1/2/18.0
 Interface ge-1/2/22.0
 RPF Nexthops: Interface ge-1/2/15.0, 1.2.94.1, 301568, 1
 Interface ge-1/2/20.0, 1.2.96.1, 301584, 65534
 Interface ge-1/2/14.0, 1.2.3.1, 301568, 1
 Interface ge-1/2/19.0, 1.3.6.1, 301584, 65534
 Attached FECs: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src: 192.168.219.11
(Active)
P2MP path type: Transit/Egress
 Output Session (label): 1.1.1.6:0 (301584) (Backup)
 Egress Nexthops: Interface ge-1/2/18.0
 Interface ge-1/2/22.0
 RPF Nexthops: Interface ge-1/2/15.0, 1.2.94.1, 301568, 1
 Interface ge-1/2/20.0, 1.2.96.1, 301584, 65534
 Interface ge-1/2/14.0, 1.2.3.1, 301568, 1
 Interface ge-1/2/19.0, 1.3.6.1, 301584, 65534
 Attached FECs: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src: 192.168.219.11
(Active)
P2MP path type: Transit/Egress
 Output Session (label): 1.1.1.6:0 (301600) (Primary)
 Egress Nexthops: Interface ge-1/2/18.0
 Interface ge-1/2/22.0
 RPF Nexthops: Interface ge-1/2/15.0, 1.2.94.1, 301616, 65534
 Interface ge-1/2/20.0, 1.2.96.1, 301600, 1
 Interface ge-1/2/14.0, 1.2.3.1, 301616, 65534
 Interface ge-1/2/19.0, 1.3.6.1, 301600, 1
 Attached FECs: P2MP root-addr 1.1.1.1, grp: 232.1.1.2, src: 192.168.219.11
(Active)
P2MP path type: Transit/Egress
 Output Session (label): 1.1.1.2:0 (301616) (Backup)
 Egress Nexthops: Interface ge-1/2/18.0
 Interface ge-1/2/22.0
 RPF Nexthops: Interface ge-1/2/15.0, 1.2.94.1, 301616, 65534
 Interface ge-1/2/20.0, 1.2.96.1, 301600, 1
 Interface ge-1/2/14.0, 1.2.3.1, 301616, 65534
 Interface ge-1/2/19.0, 1.3.6.1, 301600, 1
 Attached FECs: P2MP root-addr 1.1.1.1, grp: 232.1.1.2, src: 192.168.219.11
(Active)

```

**Meaning** The output shows primary and backup sessions, and RPF next hops.

### Checking the LDP Point-to-Multipoint Traffic Statistics

**Purpose** Make sure that both primary and backup statistics are listed.

**Action** user@R3> `show ldp traffic-statistics p2mp`

P2MP FEC Statistics:

| FEC(root_addr:lsp_id/grp,src)                                 | Nexthop | Packets | Bytes |
|---------------------------------------------------------------|---------|---------|-------|
| Shared                                                        |         |         |       |
| 1.1.1.1:232.1.1.1,192.168.219.11, Label: 301568               | 1.3.8.2 | 0       | 0     |
| No                                                            | 1.3.4.2 | 0       | 0     |
| No                                                            |         |         |       |
| 1.1.1.1:232.1.1.1,192.168.219.11, Label: 301584, Backup route | 1.3.4.2 | 0       | 0     |
| No                                                            | 1.3.8.2 | 0       | 0     |
| No                                                            |         |         |       |
| 1.1.1.1:232.1.1.2,192.168.219.11, Label: 301600               | 1.3.8.2 | 0       | 0     |
| No                                                            | 1.3.4.2 | 0       | 0     |
| No                                                            |         |         |       |
| 1.1.1.1:232.1.1.2,192.168.219.11, Label: 301616, Backup route | 1.3.4.2 | 0       | 0     |
| No                                                            | 1.3.8.2 | 0       | 0     |
| No                                                            |         |         |       |

**Meaning** The output shows both primary and backup routes with the labels.

- Related Documentation**
- [Understanding Multicast-Only Fast Reroute on page 566](#)
  - [Configuring Multicast-Only Fast Reroute on page 573](#)
  - [Example: Configuring Multicast-Only Fast Reroute in a PIM Domain](#)

## Example: Configuring LDP Downstream on Demand

This example shows how to configure LDP downstream on demand. LDP is commonly configured using downstream unsolicited advertisement mode, meaning label advertisements for all routes are received from all LDP peers. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between the access and aggregation networks and to reduce the processing requirements for the control plane.

Downstream nodes could potentially receive tens of thousands of label bindings from upstream aggregation nodes. Instead of learning and storing all label bindings for all possible loopback addresses within the entire MPLS network, the downstream aggregation node can be configured using LDP downstream on demand to only request the label bindings for the FECs corresponding to the loopback addresses of those egress nodes on which it has services configured.

- [Requirements on page 593](#)
- [Overview on page 593](#)

- [Configuration on page 593](#)
- [Verification on page 596](#)

## Requirements

This example uses the following hardware and software components:

- M Series router
- Junos OS 12.2

## Overview

You can enable LDP downstream on demand label advertisement for an LDP session by including the **downstream-on-demand** statement at the **[edit protocols ldp session]** hierarchy level. If you have configured downstream on demand, the Juniper Networks router advertises the downstream on demand request to its peer routers. For a downstream on demand session to be established between two routers, both have to advertise downstream on demand mode during LDP session establishment. If one router advertises downstream unsolicited mode and the other advertises downstream on demand, downstream unsolicited mode is used.

## Configuration

### Configuring LDP Downstream on Demand

#### Step-by-Step Procedure

To configure a LDP downstream on demand policy and then configure that policy and enable LDP downstream on demand on the LDP session:

1. Configure the downstream on demand policy (DOD-Request-Loopbacks in this example).

This policy causes the router to forward label request messages only to the FECs that are matched by the DOD-Request-Loopbacks policy.

```
[edit policy-options]
user@host# set prefix-list Request-Loopbacks 10.1.1.1/32
user@host# set prefix-list Request-Loopbacks 10.1.1.2/32
user@host# set prefix-list Request-Loopbacks 10.1.1.3/32
user@host# set prefix-list Request-Loopbacks 10.1.1.4/32
user@host# set policy-statement DOD-Request-Loopbacks term 1 from prefix-list
Request-Loopbacks
user@host# set policy-statement DOD-Request-Loopbacks term 1 then accept
```

2. Specify the DOD-Request-Loopbacks policy using the **dod-request-policy** statement at the **[edit protocols ldp]** hierarchy level.

The policy specified with the **dod-request-policy** statement is used to identify the prefixes to send label request messages. This policy is similar to an egress policy or an import policy. When processing routes from the inet.0 routing table, the Junos OS software checks for routes matching the **DOD-Request-Loopbacks** policy (in this example). If the route matches the policy and the LDP session is negotiated with DOD advertisement mode, label request messages are sent to the corresponding downstream LDP session.

```
[edit protocols ldp]
user@host# set dod-request-policy DOD-Request-Loopbacks
```

3. Include the **downstream-on-demand** statement in the configuration for the LDP session to enable downstream on demand distribution mode.

```
[edit protocols ldp]
user@host# set session 1.1.1.1 downstream-on-demand
```

### Distributing LDP Downstream on Demand Routes into Labeled BGP

**Step-by-Step Procedure** To distribute LDP downstream on demand routes into labeled BGP, use a BGP export policy.

1. Configure the LDP route policy (**redistribute\_ldp** in this example).

```
[edit policy-options]
user@host# set policy-statement redistribute_ldp term 1 from protocol ldp
user@host# set policy-statement redistribute_ldp term 1 from tag 1000
user@host# set policy-statement redistribute_ldp term 1 then accept
```

2. Include the LDP route policy, **redistribute\_ldp** in the BGP configuration (as a part of the BGP group configuration **ebgp-to-abr** in this example).

BGP forwards the LDP routes based on the **redistribute\_ldp** policy to the remote PE router

```
[edit protocols bgp]
user@host# set group ebgp-to-abr type external
user@host# set group ebgp-to-abr local-address 192.168.0.1
user@host# set group ebgp-to-abr peer-as 65319
user@host# set group ebgp-to-abr local-as 65320
user@host# set group ebgp-to-abr neighbor 192.168.6.1 family inet unicast
user@host# set group ebgp-to-abr neighbor 192.168.6.1 family inet labeled-unicast
rib inet.3
user@host# set group ebgp-to-abr neighbor 192.168.6.1 export redistribute_ldp
```

**Step-by-Step Procedure** To restrict label propagation to other routers configured in downstream unsolicited mode (instead of downstream on demand), configure the following policies:

1. Configure the **dod-routes** policy to accept routes from LDP.

```
user@host# set policy-options policy-statement dod-routes term 1 from protocol ldp
user@host# set policy-options policy-statement dod-routes term 1 from tag 1145307136
user@host# set policy-options policy-statement dod-routes term 1 then accept
```

2. Configure the **do-not-propagate-du-sessions** policy to not forward routes to neighbors 1.1.1.1, 2.2.2.2, and 3.3.3.3.

```
user@host# set policy-options policy-statement do-not-propagate-du-sessions term 1 to neighbor 1.1.1.1
user@host# set policy-options policy-statement do-not-propagate-du-sessions term 1 to neighbor 2.2.2.2
user@host# set policy-options policy-statement do-not-propagate-du-sessions term 1 to neighbor 3.3.3.3
```



```
user@host# set policy-options policy-statement do-not-propagate-du-sessions
term 1 then reject
```

3. Configure the **filter-dod-on-du-sessions** policy to prevent the routes examined by the **dod-routes** policy from being forwarded to the neighboring routers defined in the **do-not-propagate-du-sessions** policy.

```
user@host# set policy-options policy-statement filter-dod-routes-on-du-sessions
term 1 from policy dod-routes
user@host# set policy-options policy-statement filter-dod-routes-on-du-sessions
term 1 to policy do-not-propagate-du-sessions
```

4. Specify the **filter-dod-routes-on-du-session** policy as the export policy for BGP group **ebgp-to-abr**.

```
[edit protocols bgp]
user@host# set group ebgp-to-abr neighbor 192.168.6.2 export
filter-dod-routes-on-du-sessions
```

**Results** From configuration mode, confirm your configuration by entering the **show policy-options** and **show protocols ldp** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host#
show policy-options
prefix-list Request-Loopbacks {
 10.1.1.1/32;
 10.1.1.2/32;
 10.1.1.3/32;
 10.1.1.4/32;
}
policy-statement DOD-Request-Loopbacks {
 term 1 {
 from {
 prefix-list Request-Loopbacks;
 }
 then accept;
 }
}
policy-statement redistribute_ldp {
 term 1 {
 from {
 protocol ldp;
 tag 1000;
 }
 then accept;
 }
}

user@host#
show protocols ldp
dod-request-policy DOD-Request-Loopbacks;
session 1.1.1.1 {
 downstream-on-demand;
}

user@host#
```

```
show protocols bgp
group ebgp-to-abr {
 type external;
 local-address 192.168.0.1;
 peer-as 65319;
 local-as 65320;
 neighbor 192.168.6.1 {
 family inet {
 unicast;
 labeled-unicast {
 rib {
 inet.3;
 }
 }
 }
 }
 export redistribute_ldp;
}
```

## Verification

### Verifying Label Advertisement Mode

---

**Purpose** Confirm that the configuration is working properly.

Use the **show ldp session** command to verify the status of the label advertisement mode for the LDP session.

**Action** Issue the `show ldp session` and `show ldp session detail` commands:

- The following command output for the `show ldp session` command indicates that the **Adv. Mode** (label advertisement mode) is **DOD** (meaning the LDP downstream on demand session is operational):

```
user@host> show ldp session
 Address State Connection Hold time Adv. Mode
1.1.1.2 Operational Open 22 DOD
```

- The following command output for the `show ldp session detail` command indicates that the **Local Label Advertisement mode** is **Downstream unsolicited**, the default value (meaning downstream on demand is not configured on the local session). Conversely, the **Remote Label Advertisement mode** and the **Negotiated Label Advertisement mode** both indicate that **Downstream on demand** is configured on the remote session

```
user@host> show ldp session detail
Address: 1.1.1.2, State: Operational, Connection: Open, Hold time: 24
Session ID: 1.1.1.1:0--1.1.1.2:0
Next keepalive in 4 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: configured-tunneled
Keepalive interval: 10, Connect retry interval: 1
Local address: 1.1.1.1, Remote address: 1.1.1.2
Up for 17:54:52
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: disabled, Helper mode: enabled,
Remote - Restart: disabled, Helper mode: enabled
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream on demand
Negotiated Label Advertisement mode: Downstream on demand
Nonstop routing state: Not in sync
Next-hop addresses received:
 1.1.1.2
```

## Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs

- [Understanding Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs on page 598](#)
- [Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs on page 607](#)

## Understanding Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs

The Multipoint Label Distribution Protocol (M-LDP) for point-to-multipoint label-switched paths (LSPs) with in-band signaling is useful in a deployment with an existing IP/MPLS backbone, in which you need to carry multicast traffic, for IPTV for example.

For years, the most widely used solution for transporting multicast traffic has been to use native IP multicast in the service provider core with multipoint IP tunneling to isolate customer traffic. A multicast routing protocol, usually Protocol Independent Multicast (PIM), is deployed to set up the forwarding paths. IP multicast routing is used for forwarding, using PIM signaling in the core. For this model to work, the core network has to be multicast enabled. This allows for effective and stable deployments even in inter-autonomous system (AS) scenarios.

However, in an existing IP/MPLS network, deploying PIM might not be the first choice. Some service providers are interested in replacing IP tunneling with MPLS label encapsulation. The motivations for moving to MPLS label switching is to leverage MPLS traffic engineering and protection features and to reduce the amount of control traffic overhead in the provider core.

To do this, service providers are interested in leveraging the extension of the existing deployments to allow multicast traffic to pass through. The existing multicast extensions for IP/MPLS are point-to-multipoint extensions for RSVP-TE and point-to-multipoint and multipoint-to-multipoint extensions for LDP. These deployment scenarios are discussed in RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*. This feature overview is limited to point-to-multipoint extensions for LDP.

- [How M-LDP Works on page 599](#)
- [Terminology on page 603](#)
- [Ingress Join Translation and Pseudo Interface Handling on page 604](#)
- [Ingress Splicing on page 604](#)
- [Reverse Path Forwarding on page 604](#)
- [LSP Root Detection on page 604](#)
- [Egress Join Translation and Pseudo Interface Handling on page 604](#)
- [Egress Splicing on page 605](#)
- [Supported Functionality on page 605](#)
- [Unsupported Functionality on page 605](#)
- [LDP Functionality on page 606](#)
- [Egress LER Functionality on page 606](#)
- [Transit LSR Functionality on page 606](#)
- [Ingress LER Functionality on page 606](#)


### How M-LDP Works

- [Label Bindings in M-LDP Signaling on page 599](#)
- [M-LDP in PIM-Free MPLS Core on page 599](#)
- [M-LDP in PIM-Enabled MPLS Core on page 601](#)

#### Label Bindings in M-LDP Signaling

The multipoint extension to LDP uses point-to-multipoint and multipoint-to-multipoint forwarding equivalence class (FEC) elements (defined in RFC 5036, *LDP Specification*) along with capability advertisements, label mapping, and signaling procedures. The FEC elements include the idea of the LSP root, which is an IP address, and an “opaque” value, which is a selector that groups together the leaf nodes sharing the same opaque value. The opaque value is transparent to the intermediate nodes, but has meaning for the LSP root. Every LDP node advertises its local incoming label binding to the upstream LDP node on the shortest path to the root IP address found in the FEC. The upstream node receiving the label bindings creates its own local label and outgoing interfaces. This label allocation process might result in packet replication, if there are multiple outgoing branches. As shown in [Figure 61 on page 599](#), an LDP node merges the label bindings for the same opaque value if it finds downstream nodes sharing the same upstream node. This allows for effective building of point-to-multipoint LSPs and label conservation.

Figure 61: Label Bindings in M-LDP Signaling

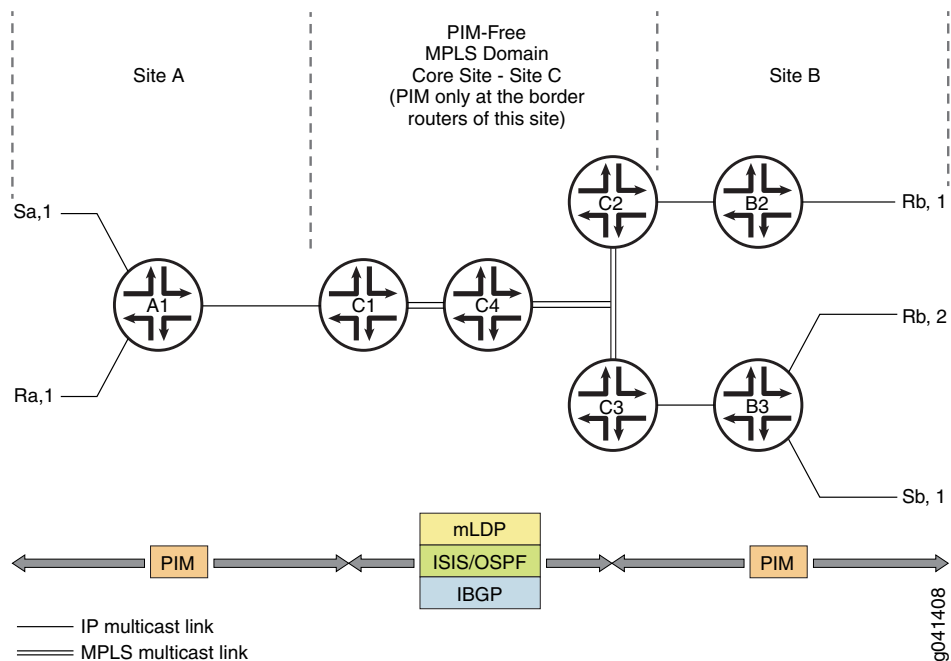


#### M-LDP in PIM-Free MPLS Core

[Figure 62 on page 600](#) shows a scaled-down deployment scenario. Two separate PIM domains are interconnected by a PIM-free core site. The border routers in this core site support PIM on the border interfaces. Further, these border routers collect and distribute the routing information from the adjacent sites to the core network. The edge routers in Site C run BGP for root-node discovery. Interior gateway protocol (IGP) routes cannot be used for ingress discovery because in most cases the forwarding next hop provided by the IGP would not provide information about the ingress device toward the source. M-LDP inband signaling has a one-to-one mapping between the point-to-multipoint LSP and the (S,G) flow. With in-band signaling, PIM messages are directly translated into M-LDP FEC bindings. In contrast, out-of-band signaling is based on manual

configuration. One application for M-LDP inband signaling is to carry IPTV multicast traffic in an MPLS backbone.

Figure 62: Sample M-LDP Topology in PIM-Free MPLS Core



### Configuration

The configuration statement **mldp-inband-signalling** on the label-edge router (LER) enables PIM to use M-LDP in-band signaling for the upstream neighbors when the LER does not detect a PIM upstream neighbor. Static configuration of the MPLS LSP root is included in the PIM configuration, using policy. This is needed when IBGP is not available in the core site or to override IBGP-based LSP root detection.

For example:

```
protocols {
 pim {
 mldp-inband-signalling {
 policy lsp-mapping-policy-example;
 }
 }
}

policy-options {
 policy-statement lsp-mapping-policy-example {
 term channel1 {
 from {
 source-address-filter ip-prefix</prefix-length>; #policy filter for channel1
 }
 then {
 p2mp-lsp-root {
 # Statically configured ingress address of edge
 # used by channel1
 }
 }
 }
 }
}
```

```


 address ip-address;
 }
 accept;
}
}
}
}

```

### M-LDP in PIM-Enabled MPLS Core

In order to migrate existing IPTV services from native IP multicast to MPLS multicast, customers need to smoothly transition from PIM to M-LDP point-to-multipoint LSPs with minimal outage. [Figure 63 on page 601](#) shows a similar M-LDP topology as [Figure 62 on page 600](#), but with a different scenario. The core is enabled with PIM, with one source streaming all the IPTV channels. The TV channels are sent as ASM streams with each channel identified by its group address. Previously, these channels were streamed on the core as IP streams and signaled using PIM.

Figure 63: Sample M-LDP Topology in PIM-Enabled MPLS Core



804263

By configuring the **mldp-inband-signaling** in this scenario, M-LDP signaling is initiated only when there is no PIM neighbor towards the source. However, because there is always a PIM neighbor towards the source unless PIM is deactivated on the upstream interfaces of the egress PE, PIM takes precedence over M-LDP and M-LDP does not take effect.

### Configuration

To progressively migrate channel by channel to M-LDP MPLS core with few streams using M-LDP upstream and other streams using existing PIM upstream, include the **selected-mldp-egress** configuration statement along with group based filters in the policy filter for M-LDP inband signaling.



**NOTE:** The M-LDP inband signaling policy filter can include either the `source-address-filter` statement or the `route-filter` statement, or a combination of both.

For example:

```
protocols {
 pim {
 mldp-inband-signalling {
 policy lsp-mapping-policy-example;
 }
 }
}

policy-options {
 policy-statement lsp-mapping-policy-example {
 term channel1 {
 from {
 source-address-filter ip-prefix</prefix-length>; #policy filter for channel1
 }
 then {
 selected-mldp-egress;
 accept;
 }
 }
 term channel2 {
 from {
 source-address-filter ip-prefix</prefix-length>; #policy filter for channel2
 route-filter ip-prefix</prefix-length>; #policy filter on multicast group address
 }
 then {
 selected-mldp-egress;
 p2mp-lsp-root {
 # Statically configured ingress address of edge
 # used by channel2
 address ip-address;
 }
 accept;
 }
 }
 term channel3 {
 from {
 route-filter ip-prefix</prefix-length>; #policy filter on multicast group address
 }
 then {
 selected-mldp-egress;
 accept;
 }
 }
 }
}
```



**NOTE:**

Some of the limitations of the above configuration are as follows:

- The `selected-mldp-egress` statement should be configured only on the LER. Configuring the `selected-mldp-egress` statement on non-egress PIM routers can cause path setup failures.
- When policy changes are made to switch traffic from PIM upstream to M-LDP upstream and vice-versa, packet loss can be expected as break-and-make mechanism is performed at the control plane.

### Terminology

The following terms are important for an understanding of M-LDP in-band signaling for multicast traffic.

**Point-to-point LSP**—An LSP that has one ingress label-switched router (LSR) and one egress LSR.

**Multipoint LSP**—Either a point-to-multipoint or a multipoint-to-multipoint LSP.

**Point-to-multipoint LSP**—An LSP that has one ingress LSR and one or more egress LSRs.

**Multipoint-to-point LSP**—An LSP that has one or more ingress LSRs and one unique egress LSR.

**Multipoint-to-multipoint LSP**—An LSP that connects a set of nodes, such that traffic sent by any node in the LSP is delivered to all others.

**Ingress LSR**—An ingress LSR for a particular LSP is an LSR that can send a data packet along the LSP. Multipoint-to-multipoint LSPs can have multiple ingress LSRs. Point-to-multipoint LSPs have only one, and that node is often referred to as the root node.

**Egress LSR**—An egress LSR for a particular LSP is an LSR that can remove a data packet from that LSP for further processing. Point-to-point and multipoint-to-point LSPs have only a single egress node. Point-to-multipoint and multipoint-to-multipoint LSPs can have multiple egress nodes.

**Transit LSR**—An LSR that has reachability to the root of the multipoint LSP through a directly connected upstream LSR and one or more directly connected downstream LSRs.

**Bud LSR**—An LSR that is an egress but also has one or more directly connected downstream LSRs.

**Leaf node**—Either an egress or bud LSR in the context of a point-to-multipoint LSP. In the context of a multipoint-to-multipoint LSP, an LSR is both ingress and egress for the same multipoint-to-multipoint LSP and can also be a bud LSR.

### Ingress Join Translation and Pseudo Interface Handling

---

At the ingress LER, LDP notifies PIM about the (S,G) messages that are received over the in-band signaling. PIM associates each (S,G) message with a pseudo interface. Subsequently, a shortest-path-tree (SPT) join message is initiated toward the source. PIM treats this as a new type of local receiver. When the LSP is torn down, PIM removes this local receiver based on notification from LDP.

### Ingress Splicing

---

LDP provides PIM with a next hop to be associated with each (S,G) entry. PIM installs a PIM (S,G) multicast route with the LDP next hop and other PIM receivers. The next hop is a composite next hop of local receivers + the list of PIM downstream neighbors + a sub-level next hop for the LDP tunnel.

### Reverse Path Forwarding

---

PIM's reverse-path-forwarding (RPF) calculation is performed at the egress node.

PIM performs M-LDP in-band signaling when all of the following conditions are true:

- There are no PIM neighbors toward the source.
- The M-LDP in-band signaling statement is configured.
- The next hop is learned through BGP, or is present in the static mapping (specified in an M-LDP in-band signaling policy).

Otherwise, if LSP root detection fails, PIM retains the (S,G) entry with an RPF state of unresolved.

PIM RPF registers this source address each time unicast routing information changes. Therefore, if the route toward the source changes, the RPF recalculation recurs. BGP protocol next hops toward the source too are monitored for changes in the LSP root. Such changes might cause traffic disruption for short durations.

### LSP Root Detection

---

If the RPF operation detects the need for M-LDP in-band signaling upstream, the LSP root (ingress) is detected. This root is a parameter for LDP LSP signaling.

The root node is detected as follows:

1. If the existing static configuration specifies the source address, the root is taken as given in configuration.
2. A lookup is performed in the unicast routing table. If the source address is found, the protocol next hop toward the source is used as the LSP root.

### Egress Join Translation and Pseudo Interface Handling

---

At the egress LER, PIM notifies LDP of the (S,G) message to be signaled along with the LSP root. PIM creates a pseudo interface as the upstream interface for this (S,G) message. When an (S,G) prune message is received, this association is removed.

### Egress Splicing

---

At the egress node of the core network, where the (S,G) join message from the downstream site is received, this join message is translated to M-LDP in-band signaling parameters and LDP is notified. Further, LSP teardown occurs when the (S,G) entry is lost, when the LSP root changes, or when the (S,G) entry is reachable over a PIM neighbor.

### Supported Functionality

---

For M-LDP in-band signaling, Junos OS supports the following functionality:

- Egress splicing of the PIM next hop with the LDP route
- Ingress splicing of the PIM route with the LDP next hop
- Translation of PIM join messages to LDP point-to-multipoint LSP setup parameters
- Translation of M-LDP in-band LSP parameters to set up PIM join messages
- Statically configured and BGP protocol next hop-based LSP root detection
- PIM (S,G) states in the PIM source-specific multicast (SSM) and anysource multicast (ASM) ranges
- Configuration statements on ingress and egress LERs to enable them to act as edge routers
- IGMP join messages on LERs
- Carrying IPv6 source and group address as opaque information toward an IPv4 root node
- Static configuration to map an IPv6 (S,G) to an IPv4 root address

### Unsupported Functionality

---

For M-LDP in-band signaling, Junos OS does *not* support the following functionality:

- Full support for PIM ASM
- The **mpls lsp point-to-multipoint ping** command with an (S,G) option
- Nonstop active routing (NSR)
- Make-before-break (MBB) for PIM
- IPv6 LSP root addresses (LDP does not support IPv6 LSPs.)
- Neighbor relationship between PIM speakers that are not directly connected
- Graceful restart
- PIM dense mode
- PIM bidirectional mode

### LDP Functionality

---

The PIM (S,G) information is carried as M-LDP opaque type-length-value (TLV) encodings. The point-to-multipoint FEC element consists of the root-node address. In the case of next-generation multicast VPNs (NGEN MVPNs), the point-to-multipoint LSP is identified by the root node address and the LSP ID.

### Egress LER Functionality

---

On the egress LER, PIM triggers LDP with the following information to create a point-to-multipoint LSP:

- Root node
- (S,G)
- Next hop

PIM finds the root node based on the source of the multicast tree. If the root address is configured for this (S,G) entry, the configured address is used as the point-to-multipoint LSP root. Otherwise, the routing table is used to look up the route to the source. If the route to the source of the multicast tree is a BGP-learned route, PIM retrieves the BGP next hop address and uses it as the root node for the point-to-multipoint LSP.

LDP finds the upstream node based on the root node, allocates a label, and sends the label mapping to the upstream node. LDP does not use penultimate hop popping (PHP) for in-band M-LDP signaling.

If the root addresses for the source of the multicast tree changes, PIM deletes the point-to-multipoint LSP and triggers LDP to create a new point-to-multipoint LSP. When this happens, the outgoing interface list becomes NULL, PIM triggers LDP to delete the point-to-multipoint LSP, and LDP sends a label withdraw message to the upstream node.

### Transit LSR Functionality

---

The transit LSR advertises a label to the upstream LSR toward the source of the point-to-multipoint FEC and installs the necessary forwarding state to forward the packets. The transit LSR can be any M-LDP capable router.

### Ingress LER Functionality

---

On the ingress LER, LDP provides the following information to PIM upon receiving the label mapping:

- (S,G)
- Flood next hop

Then PIM installs the forwarding state. If the new branches are added or deleted, the flood next hop is updated accordingly. If all branches are deleted due to a label being withdrawn, LDP sends updated information to PIM. If there are multiple links between

the upstream and downstream neighbors, the point-to-multipoint LSP is not load balanced.

## Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs

This example shows how to configure multipoint LDP (M-LDP) in-band signaling for multicast traffic, as an extension to the Protocol Independent Multicast (PIM) protocol or as a substitute for PIM.

- [Requirements on page 607](#)
- [Overview on page 607](#)
- [Configuration on page 608](#)
- [Verification on page 617](#)

### Requirements

This example can be configured using the following hardware and software components:

- Junos OS Release 13.2 or later
- MX Series 3D Universal Edge Routers or M Series Multiservice Edge Routers for the Provider Edge (PE) Routers
- PTX Series Packet Transport Routers acting as transit label-switched routers
- T Series Core Routers for the Core Routers



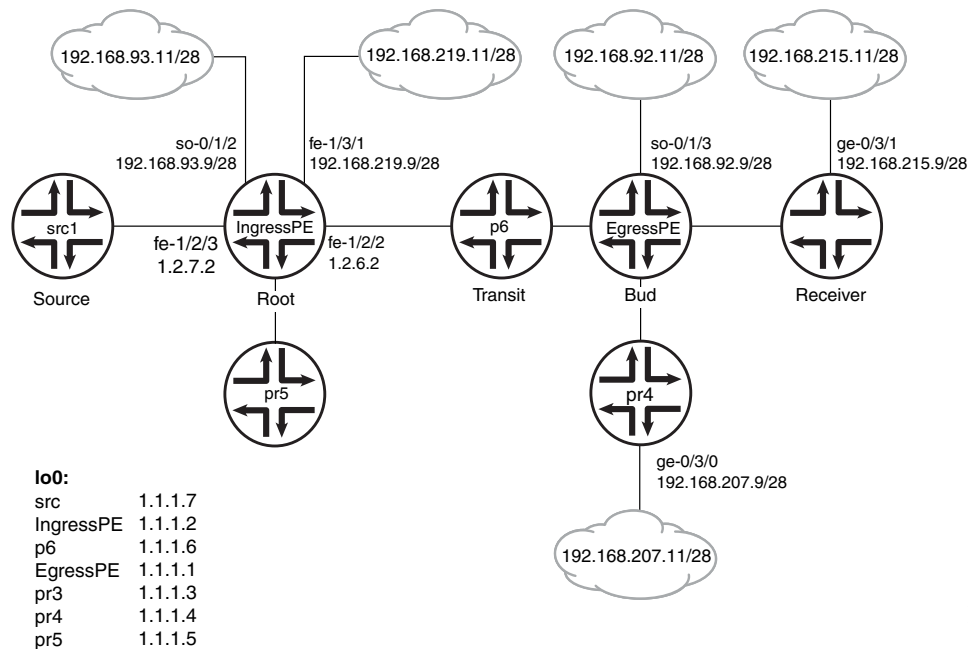
**NOTE:** The PE routers could also be T Series Core Routers but that is not typical. Depending on your scaling requirements, the core routers could also be MX Series 3D Universal Edge Routers or M Series Multiservice Edge Routers. The Customer Edge (CE) devices could be other routers or switches from Juniper Networks or another vendor.

No special configuration beyond device initialization is required before configuring this example.

### Overview

"CLI Quick Configuration" on page 608 shows the configuration for all of the devices in Figure 64 on page 608. The section "Step-by-Step Procedure" on page 611 describes the steps on Device EgressPE.

Figure 64: M-LDP In-Band Signaling for Point-to-Multipoint LSPs Example Topology



### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

#### Device src1

```
set logical-systems src1 interfaces fe-1/2/0 unit 0 family inet address 1.2.7.7/24
set logical-systems src1 interfaces lo0 unit 0 family inet address 1.1.1.7/32
set logical-systems src1 protocols ospf area 0.0.0.0 interface all
```

#### Device IngressPE

```
set interfaces so-0/1/2 unit 0 family inet address 192.168.93.9/28
set interfaces fe-1/2/0 unit 0 family inet address 1.2.3.2/24
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 1.2.5.2/24
set interfaces fe-1/2/2 unit 0 family inet address 1.2.6.2/24
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces fe-1/2/3 unit 0 family inet address 1.2.7.2/24
set interfaces fe-1/3/1 unit 0 family inet address 192.168.219.9/28
set interfaces lo0 unit 0 family inet address 1.1.1.2/32
set protocols igmp interface fe-1/2/1.0 version 3
set protocols igmp interface fe-1/2/1.0 static group 232.1.1.1 source 192.168.219.11
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.2
set protocols bgp group ibgp family inet any
set protocols bgp group ibgp family inet-vpn any
set protocols bgp group ibgp neighbor 1.1.1.3
set protocols bgp group ibgp neighbor 1.1.1.4
set protocols bgp group ibgp neighbor 1.1.1.1
```

```

set protocols ospf area 0.0.0.0 interface all
set protocols ldp interface fe-1/2/0.0
set protocols ldp interface fe-1/2/2.0
set protocols ldp interface lo0.0
set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim rp static address 1.1.1.5
set protocols pim interface fe-1/3/1.0
set protocols pim interface lo0.0
set protocols pim interface fe-1/2/0.21
set protocols pim interface fe-1/2/3.0
set protocols pim interface fe-1/2/1.0
set protocols pim interface so-0/1/2.0
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.0.0/24 orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term A from source-address-filter
 1.1.1.7/32 orlonger
set policy-options policy-statement mldppim-ex term A from source-address-filter
 1.2.7.0/24 orlonger
set policy-options policy-statement mldppim-ex term A then accept
set routing-options autonomous-system 64510

```

#### Device EgressPE

```

set interfaces so-0/1/3 unit 0 point-to-point
set interfaces so-0/1/3 unit 0 family inet address 192.168.92.9/28
set interfaces fe-1/2/0 unit 0 family inet address 1.1.3.1/24
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 1.1.4.1/24
set interfaces fe-1/2/2 unit 0 family inet address 1.1.6.1/24
set interfaces fe-1/2/2 unit 0 family mpls
set interfaces fe-1/3/0 unit 0 family inet address 192.168.209.9/28
set interfaces lo0 unit 0 family inet address 1.1.1.1/32
set routing-options autonomous-system 64510
set protocols igmp interface fe-1/3/0.0 version 3
set protocols igmp interface fe-1/3/0.0 static group 232.1.1.1 group-count 3
set protocols igmp interface fe-1/3/0.0 static group 232.1.1.1 source 192.168.219.11
set protocols igmp interface fe-1/3/0.0 static group 227.1.1.1
set protocols igmp interface so-0/1/3.0 version 3
set protocols igmp interface so-0/1/3.0 static group 232.1.1.1 group-count 2
set protocols igmp interface so-0/1/3.0 static group 232.1.1.1 source 192.168.219.11
set protocols igmp interface so-0/1/3.0 static group 232.2.2.2 source 1.2.7.7
set protocols mpls interface fe-1/2/0.0
set protocols mpls interface fe-1/2/2.0
set protocols bgp group ibgp type internal
set protocols bgp group ibgp local-address 1.1.1.1
set protocols bgp group ibgp family inet any
set protocols bgp group ibgp neighbor 1.1.1.2
set protocols msdp local-address 1.1.1.1
set protocols msdp peer 1.1.1.5
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface fe-1/2/0.0
set protocols ldp interface fe-1/2/2.0
set protocols ldp interface lo0.0

```

```

set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim rp local address 1.1.1.1
set protocols pim rp local group-ranges 227.0.0.0/8
set protocols pim rp static address 1.1.1.4
set protocols pim rp static address 1.2.7.7 group-ranges 226.0.0.0/8
set protocols pim interface lo0.0
set protocols pim interface fe-1/3/0.0
set protocols pim interface fe-1/2/0.0
set protocols pim interface fe-1/2/1.0
set protocols pim interface so-0/1/3.0
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.0.0/24 orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address
 1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term A from source-address-filter
 1.2.7.0/24 orlonger
set policy-options policy-statement mldppim-ex term A then accept

```

**Device p6**

```

set interfaces fe-1/2/0 unit 0 family inet address 1.1.6.6/24
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 1.2.6.6/24
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.6/32
set interfaces lo0 unit 0 family mpls
set protocols ospf area 0.0.0.0 interface all
set protocols ldp interface fe-1/2/0.0
set protocols ldp interface fe-1/2/1.0
set protocols ldp interface lo0.0
set protocols ldp p2mp

```

**Device pr3**

```

set interfaces ge-0/3/1 unit 0 family inet address 192.168.215.9/28
set interfaces fe-1/2/0 unit 0 family inet address 1.1.3.3/24
set interfaces fe-1/2/0 unit 0 family mpls
set interfaces fe-1/2/1 unit 0 family inet address 1.2.3.3/24
set interfaces fe-1/2/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 1.1.1.3/32
set protocols igmp interface ge-0/3/1.0 version 3
set protocols igmp interface ge-0/3/1.0 static group 232.1.1.2 source 192.168.219.11
set protocols igmp interface ge-0/3/1.0 static group 232.2.2.2 source 1.2.7.7
set protocols bgp group ibgp local-address 1.1.1.3
set protocols bgp group ibgp type internal
set protocols bgp group ibgp neighbor 1.1.1.2
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fe-1/2/1.0 metric 2
set protocols ldp interface fe-1/2/0.0
set protocols ldp interface fe-1/2/1.0
set protocols ldp interface lo0.0
set protocols ldp p2mp
set protocols pim mldp-inband-signalling policy mldppim-ex
set protocols pim interface fe-0/3/1.0
set protocols pim interface lo0.0

```



```

set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.0.0/24 orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address
 1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.0.0/24 orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
 192.168.219.11/32 orlonger
set policy-options policy-statement mldppim-ex term B from source-address-filter
 1.2.7.7/32 orlonger
set policy-options policy-statement mldppim-ex term B then p2mp-lsp-root address
 1.1.1.2
set policy-options policy-statement mldppim-ex term B then accept
set routing-options autonomous-system 64510

```

**Device pr4**

```

set interfaces ge-0/3/0 unit 0 family inet address 192.168.207.9/28
set interfaces fe-1/2/0 unit 0 family inet address 1.1.4.4/24
set interfaces fe-1/2/0 unit 0 family iso
set interfaces lo0 unit 0 family inet address 1.1.1.4/32
set protocols igmp interface ge-0/3/0.0 version 3
set protocols igmp interface ge-0/3/0.0 static group 232.1.1.2 source 192.168.219.11
set protocols igmp interface ge-0/3/0.0 static group 225.1.1.1
set protocols bgp group ibgp local-address 1.1.1.4
set protocols bgp group ibgp type internal
set protocols bgp group ibgp neighbor 1.1.1.2
set protocols msdp local-address 1.1.1.4
set protocols msdp peer 1.1.1.5
set protocols ospf area 0.0.0.0 interface all
set protocols pim rp local address 1.1.1.4
set protocols pim interface ge-0/3/0.0
set protocols pim interface lo0.0
set protocols pim interface fe-1/2/0.0
set routing-options autonomous-system 64510

```

**Device pr5**

```

set interfaces fe-1/2/0 unit 0 family inet address 1.2.5.5/24
set interfaces lo0 unit 0 family inet address 1.1.1.5/24
set protocols igmp interface lo0.0 version 3
set protocols igmp interface lo0.0 static group 232.1.1.1 source 192.168.219.11
set protocols msdp local-address 1.1.1.5
set protocols msdp peer 1.1.1.4
set protocols msdp peer 1.1.1.1
set protocols ospf area 0.0.0.0 interface all
set protocols pim rp local address 1.1.1.5
set protocols pim interface all

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device EgressPE:

1. Configure the interfaces.

Enable MPLS on the core-facing interfaces. On the egress next hops, you do not need to enable MPLS.

```
[edit interfaces]
```

```
user@EgressPE# set fe-1/2/0 unit 0 family inet address 1.1.3.1/24
```

```
user@EgressPE# set fe-1/2/0 unit 0 family mpls
```

```
user@EgressPE# set fe-1/2/2 unit 0 family inet address 1.1.6.1/24
```

```
user@EgressPE# set fe-1/2/2 unit 0 family mpls
```

```
user@EgressPE# set so-0/1/3 unit 0 point-to-point
```

```
user@EgressPE# set so-0/1/3 unit 0 family inet address 192.168.92.9/28
```

```
user@EgressPE# set fe-1/2/1 unit 0 family inet address 1.1.4.1/24
```

```
user@EgressPE# set fe-1/3/0 unit 0 family inet address 192.168.209.9/28
```

```
user@EgressPE# set lo0 unit 0 family inet address 1.1.1.1/32
```

2. Configure IGMP on the egress interfaces.

For testing purposes, this example includes static group and source addresses.

```
[edit protocols igmp]
```

```
user@EgressPE# set interface fe-1/3/0.0 version 3
```

```
user@EgressPE# set interface fe-1/3/0.0 static group 232.1.1.1 group-count 3
```

```
user@EgressPE# set interface fe-1/3/0.0 static group 232.1.1.1 source 192.168.219.11
```

```
user@EgressPE# set interface fe-1/3/0.0 static group 227.1.1.1
```

```
user@EgressPE# set interface so-0/1/3.0 version 3
```

```
user@EgressPE# set interface so-0/1/3.0 static group 232.1.1.1 group-count 2
```

```
user@EgressPE# set interface so-0/1/3.0 static group 232.1.1.1 source 192.168.219.11
```

```
user@EgressPE# set interface so-0/1/3.0 static group 232.2.2.2 source 1.2.7.7
```

3. Configure MPLS on the core-facing interfaces.

```
[edit protocols mpls]
```

```
user@EgressPE# set interface fe-1/2/0.0
```

```
user@EgressPE# set interface fe-1/2/2.0
```

4. Configure BGP.

BGP is a policy-driven protocol, so also configure and apply any needed routing policies.

For example, you might want to export static routes into BGP.

```
[edit protocols bgp group ibgp]
```

```
user@EgressPE# set type internal
```

```
user@EgressPE# set local-address 1.1.1.1
```

```
user@EgressPE# set family inet any
```

```
user@EgressPE# set neighbor 1.1.1.2
```

5. (Optional) Configure an MSDP peer connection with Device pr5 in order to interconnect the disparate PIM domains, thus enabling redundant RPs.

```
[edit protocols msdp]
```

```
user@EgressPE# set local-address 1.1.1.1
```

```
user@EgressPE# set peer 1.1.1.5
```

6. Configure OSPF.
 

```
[edit protocols ospf area 0.0.0.0]
user@EgressPE# set interface all
user@EgressPE# set interface fxp0.0 disable
```
7. Configure LDP on the core-facing interfaces and on the loopback interface.
 

```
[edit protocols ldp]
user@EgressPE# set interface fe-1/2/0.0
user@EgressPE# set interface fe-1/2/2.0
user@EgressPE# set interface lo0.0
```
8. Enable point-to-multipoint MPLS LSPs.
 

```
[edit protocols ldp]
user@EgressPE# set p2mp
```
9. Configure PIM on the downstream interfaces.
 

```
[edit protocols pim]
user@EgressPE# set interface lo0.0
user@EgressPE# set interface fe-1/3/0.0
user@EgressPE# set interface fe-1/2/1.0
user@EgressPE# set interface so-0/1/3.0
```
10. Configure the RP settings because this device serves as the PIM rendezvous point (RP).
 

```
[edit protocols pim]
user@EgressPE# set rp local address 1.1.1.1
user@EgressPE# set rp local group-ranges 227.0.0.0/8
user@EgressPE# set rp static address 1.1.1.4
user@EgressPE# set rp static address 1.2.7.7 group-ranges 226.0.0.0/8
```
11. Enable M-LDP in-band signaling and set the associated policy.
 

```
[edit protocols pim]
user@EgressPE# set mldp-inband-signalling policy mldppim-ex
```
12. Configure the routing policy that specifies the root address for the point-to-multipoint LSP and the associated source addresses.
 

```
[edit policy-options policy-statement mldppim-ex]
user@EgressPE# set term B from source-address-filter 192.168.0.0/24 orlonger
user@EgressPE# set term B from source-address-filter 192.168.219.11/32 orlonger
user@EgressPE# set term B then p2mp-lsp-root address 1.1.1.2
user@EgressPE# set term B then accept

user@EgressPE# set term A from source-address-filter 1.2.7.0/24 orlonger
user@EgressPE# set term A then accept
```
13. Configure the autonomous system (AS) ID.
 

```
[edit routing-options]
user@EgressPE# set autonomous-system 64510
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output

does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
Device EgressPE user@EgressPE# show interfaces
so-0/1/3 {
 unit 0 {
 point-to-point;
 family inet {
 address 192.168.92.9/28;
 }
 }
}
fe-1/2/0 {
 unit 0 {
 family inet {
 address 1.1.3.1/24;
 }
 family mpls;
 }
}
fe-1/2/1 {
 unit 0 {
 family inet {
 address 1.1.4.1/24;
 }
 }
}
fe-1/2/2 {
 unit 0 {
 family inet {
 address 1.1.6.1/24;
 }
 family mpls;
 }
}
fe-1/3/0 {
 unit 0 {
 family inet {
 address 192.168.209.9/28;
 }
 }
}
lo0 {
 unit 0 {
 family inet {
 address 1.1.1.1/32;
 }
 }
}

user@EgressPE# show protocols
igmp {
 interface fe-1/3/0.0 {
 version 3;
 static {
 group 232.1.1.1 {
```

```

 group-count 3;
 source 192.168.219.11;
 }
 group 227.1.1.1;
}
}
interface so-0/1/3.0 {
 version 3;
 static {
 group 232.1.1.1 {
 group-count 2;
 source 192.168.219.11;
 }
 group 232.2.2.2 {
 source 1.2.7.7;
 }
 }
}
}
mpls {
 interface fe-1/2/0.0;
 interface fe-1/2/2.0;
}
bgp {
 group ibgp {
 type internal;
 local-address 1.1.1.1;
 family inet {
 any;
 }
 neighbor 1.1.1.2;
 }
}
msdp {
 local-address 1.1.1.1;
 peer 1.1.1.5;
}
ospf {
 area 0.0.0.0 {
 interface all;
 interface fxp0.0 {
 disable;
 }
 }
}
}
ldp {
 interface fe-1/2/0.0;
 interface fe-1/2/2.0;
 interface lo0.0;
 p2mp;
}
pim {
 mldp-inband-signalling {
 policy mldppim-ex;
 }
}
rp {

```

```
local {
 address 1.1.1.1;
 group-ranges {
 227.0.0.0/8;
 }
}
static {
 address 1.1.1.4;
 address 1.2.7.7 {
 group-ranges {
 226.0.0.0/8;
 }
 }
}
}
interface lo0.0;
interface fe-1/3/0.0;
interface fe-1/2/0.0;
interface fe-1/2/1.0;
interface so-0/1/3.0;
}

user@EgressPE# show policy-options
policy-statement mldppim-ex {
 term B {
 from {
 source-address-filter 192.168.0.0/24 orlonger;
 source-address-filter 192.168.219.11/32 orlonger;
 }
 then {
 p2mp-lsp-root {
 address 1.1.1.2;
 }
 accept;
 }
 }
 term A {
 from {
 source-address-filter 1.2.7.0/24 orlonger;
 }
 then accept;
 }
}

user@EgressPE# show routing-options
autonomous-system 64510;
```

Similarly, configure the other egress devices.

If you are done configuring the devices, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Checking the PIM Join States on page 617](#)
- [Checking the PIM Sources on page 620](#)
- [Checking the LDP Database on page 623](#)
- [Looking Up the Route Information for the MPLS Label on page 626](#)
- [Checking the LDP Traffic Statistics on page 627](#)

### Checking the PIM Join States

**Purpose** Display information about PIM join states to verify the M-LDP in-band upstream and downstream details. On the ingress device, the **show pim join extensive** command displays **Pseudo-MLDP** for the downstream interface. On the egress, the **show pim join extensive** command displays **Pseudo-MLDP** for the upstream interface.

**Action** From operational mode, enter the **show pim join extensive** command.

```
user@IngressPE> show pim join extensive
```

```
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

```
Group: 232.1.1.1
 Source: 192.168.219.11
 Flags: sparse,spt
 Upstream interface: fe-1/3/1.0
 Upstream neighbor: Direct
 Upstream state: Local Source
 Keepalive timeout:
 Uptime: 1d 23:00:12
 Downstream neighbors:
 Interface: Pseudo-MLDP
 Interface: fe-1/2/1.0
 1.2.5.2 State: Join Flags: S Timeout: Infinity
 Uptime: 1d 23:00:12 Time since last Join: 1d 23:00:12
```

```
Group: 232.1.1.2
 Source: 192.168.219.11
 Flags: sparse,spt
 Upstream interface: fe-1/3/1.0
 Upstream neighbor: Direct
 Upstream state: Local Source
 Keepalive timeout:
 Uptime: 1d 22:59:59
 Downstream neighbors:
 Interface: Pseudo-MLDP
```

```
Group: 232.1.1.3
 Source: 192.168.219.11
 Flags: sparse,spt
 Upstream interface: fe-1/3/1.0
 Upstream neighbor: Direct
 Upstream state: Local Source
 Keepalive timeout:
```

```
Uptime: 1d 22:07:31
Downstream neighbors:
 Interface: Pseudo-MLDP

Group: 232.2.2.2
Source: 1.2.7.7
Flags: sparse,spt
Upstream interface: fe-1/2/3.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 1d 22:59:59
Downstream neighbors:
 Interface: Pseudo-MLDP

user@EgressPE> show pim join extensive

Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 227.1.1.1
Source: *
RP: 1.1.1.1
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 1d 23:14:21
Downstream neighbors:
 Interface: fe-1/3/0.0
 192.168.209.9 State: Join Flags: SRW Timeout: Infinity
 Uptime: 1d 23:14:21 Time since last Join: 1d 20:12:35

Group: 232.1.1.1
Source: 192.168.219.11
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 1d 23:14:22
Downstream neighbors:
 Interface: so-0/1/3.0
 192.168.92.9 State: Join Flags: S Timeout: Infinity
 Uptime: 1d 20:12:35 Time since last Join: 1d 20:12:35
 Downstream neighbors:
 Interface: fe-1/3/0.0
 192.168.209.9 State: Join Flags: S Timeout: Infinity
 Uptime: 1d 20:12:35 Time since last Join: 1d 20:12:35

Group: 232.1.1.2
Source: 192.168.219.11
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 1d 23:14:22
Downstream neighbors:
```



```

 Interface: so-0/1/3.0
 192.168.92.9 State: Join Flags: S Timeout: Infinity
 Uptime: 1d 20:12:35 Time since last Join: 1d 20:12:35
 Downstream neighbors:
 Interface: fe-1/2/1.0
 1.1.4.4 State: Join Flags: S Timeout: 198
 Uptime: 1d 22:59:59 Time since last Join: 00:00:12
 Downstream neighbors:
 Interface: fe-1/3/0.0
 192.168.209.9 State: Join Flags: S Timeout: Infinity
 Uptime: 1d 20:12:35 Time since last Join: 1d 20:12:35

Group: 232.1.1.3
Source: 192.168.219.11
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 1d 20:12:35
Downstream neighbors:
 Interface: fe-1/3/0.0
 192.168.209.9 State: Join Flags: S Timeout: Infinity
 Uptime: 1d 20:12:35 Time since last Join: 1d 20:12:35

Group: 232.2.2.2
Source: 1.2.7.7
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 1d 20:12:35
Downstream neighbors:
 Interface: so-0/1/3.0
 192.168.92.9 State: Join Flags: S Timeout: Infinity
 Uptime: 1d 20:12:35 Time since last Join: 1d 20:12:35

user@pr3> show pim join extensive

Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.2
Source: 192.168.219.11
Flags: sparse,spt
Upstream protocol: MLDP
Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 1d 20:14:40
Downstream neighbors:
 Interface: Pseudo-GMP
 ge-0/3/1.0

Group: 232.2.2.2
Source: 1.2.7.7
Flags: sparse,spt
Upstream protocol: MLDP

```

```

Upstream interface: Pseudo MLDP
Upstream neighbor: MLDP LSP root <1.1.1.2>
Upstream state: Join to Source
Keepalive timeout:
Uptime: 1d 20:14:40
Downstream neighbors:
 Interface: Pseudo-GMP
 ge-0/3/1.0

```

```

user@pr4> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

```

Group: 225.1.1.1
Source: *
RP: 1.1.1.4
Flags: sparse,rptree,wildcard
Upstream interface: Local
Upstream neighbor: Local
Upstream state: Local RP
Uptime: 1d 23:13:43
Downstream neighbors:
 Interface: ge-0/3/0.0
 192.168.207.9 State: Join Flags: SRW Timeout: Infinity
 Uptime: 1d 23:13:43 Time since last Join: 1d 23:13:43

```

```

Group: 232.1.1.2
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/2/0.0
Upstream neighbor: 1.1.4.1
Upstream state: Local RP, Join to Source
Keepalive timeout: 0
Uptime: 1d 23:13:43
Downstream neighbors:
 Interface: ge-0/3/0.0
 192.168.207.9 State: Join Flags: S Timeout: Infinity
 Uptime: 1d 23:13:43 Time since last Join: 1d 23:13:43

```

```

user@pr5> show pim join extensive
ge-0/3/1.0

```

```

Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

```

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

### Checking the PIM Sources

**Purpose** Verify that the PIM sources have the expected M-LDP in-band upstream and downstream details.

**Action** From operational mode, enter the **show pim source** command.

```
user@IngressPE> show pim source
```

```
Instance: PIM.master Family: INET
```

```

Source 1.1.1.1
 Prefix 1.1.1.1/32
 Upstream interface Local
 Upstream neighbor Local

Source 1.2.7.7
 Prefix 1.2.7.0/24
 Upstream protocol MLDP
 Upstream interface Pseudo MLDP
 Upstream neighbor MLDP LSP root <1.1.1.2>

Source 192.168.219.11
 Prefix 192.168.219.0/28
 Upstream protocol MLDP
 Upstream interface Pseudo MLDP
 Upstream neighbor MLDP LSP root <1.1.1.2>

user@EgressPE> show pim source
Instance: PIM.master Family: INET

Source 1.2.7.7
 Prefix 1.2.7.0/24
 Upstream interface fe-1/2/3.0
 Upstream neighbor 1.2.7.2

Source 1.2.7.7
 Prefix 1.2.7.0/24
 Upstream interface fe-1/2/3.0
 Upstream neighbor Direct

Source 192.168.219.11
 Prefix 192.168.219.0/28
 Upstream interface fe-1/3/1.0
 Upstream neighbor 192.168.219.9

Source 192.168.219.11
 Prefix 192.168.219.0/28
 Upstream interface fe-1/3/1.0
 Upstream neighbor Direct

user@pr3> show pim source

Instance: PIM.master Family: INET

Source 1.2.7.7
 Prefix 1.2.7.0/24
 Upstream protocol MLDP
 Upstream interface Pseudo MLDP
 Upstream neighbor MLDP LSP root <1.1.1.2>

Source 192.168.219.11
 Prefix 192.168.219.0/28
 Upstream protocol MLDP
 Upstream interface Pseudo MLDP
 Upstream neighbor MLDP LSP root <1.1.1.2>

user@pr4> show pim source
Instance: PIM.master Family: INET

Source 1.1.1.4

```

```
Prefix 1.1.1.4/32
Upstream interface Local
Upstream neighbor Local

Source 192.168.219.11
Prefix 192.168.219.0/28
Upstream interface fe-1/2/0.0
Upstream neighbor 1.1.4.1
```

*Checking the LDP Database*

**Purpose** Make sure that the `show ldp database` command displays the expected root-to-(S,G) bindings.

```

Action user@IngressPE> show ldp database

Input label database, 10.255.2.227:0--1.1.1.3:0
 Label Prefix
 300096 1.1.1.2/32
 3 1.1.1.3/32
 299856 1.1.1.6/32
 299776 10.255.2.227/32

Output label database, 10.255.2.227:0--1.1.1.3:0
 Label Prefix
 300144 1.1.1.2/32
 299776 1.1.1.3/32
 299856 1.1.1.6/32
 3 10.255.2.227/32

Input label database, 10.255.2.227:0--1.1.1.6:0
 Label Prefix
 299936 1.1.1.2/32
 299792 1.1.1.3/32
 3 1.1.1.6/32
 299776 10.255.2.227/32

Output label database, 10.255.2.227:0--1.1.1.6:0
 Label Prefix
 300144 1.1.1.2/32
 299776 1.1.1.3/32
 299856 1.1.1.6/32
 3 10.255.2.227/32
 300432 P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
 300288 P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
 300160 P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
 300480 P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11

user@EgressPE> show ldp database

Input label database, 1.1.1.2:0--1.1.1.3:0
 Label Prefix
 300096 1.1.1.2/32
 3 1.1.1.3/32
 299856 1.1.1.6/32
 299776 10.255.2.227/32
 300144 P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
 300128 P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11

Output label database, 1.1.1.2:0--1.1.1.3:0
 Label Prefix
 3 1.1.1.2/32
 299776 1.1.1.3/32
 299808 1.1.1.6/32
 299792 10.255.2.227/32

Input label database, 1.1.1.2:0--1.1.1.6:0
 Label Prefix
 299936 1.1.1.2/32
 299792 1.1.1.3/32
 3 1.1.1.6/32
 299776 10.255.2.227/32
 300128 P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
 299984 P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
 299952 P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11

```

```

300176 P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
300192 P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7

```

```
Output label database, 1.1.1.2:0--1.1.1.6:0
```

```

Label Prefix
3 1.1.1.2/32
299776 1.1.1.3/32
299808 1.1.1.6/32
299792 10.255.2.227/32

```

```
logical-system: default
```

```
Input label database, 10.255.2.227:0--1.1.1.3:0
```

```

Label Prefix
300096 1.1.1.2/32
3 1.1.1.3/32
299856 1.1.1.6/32
299776 10.255.2.227/32

```

```
Output label database, 10.255.2.227:0--1.1.1.3:0
```

```

Label Prefix
300144 1.1.1.2/32
299776 1.1.1.3/32
299856 1.1.1.6/32
3 10.255.2.227/32

```

```
Input label database, 10.255.2.227:0--1.1.1.6:0
```

```

Label Prefix
299936 1.1.1.2/32
299792 1.1.1.3/32
3 1.1.1.6/32
299776 10.255.2.227/32

```

```
Output label database, 10.255.2.227:0--1.1.1.6:0
```

```

Label Prefix
300144 1.1.1.2/32
299776 1.1.1.3/32
299856 1.1.1.6/32
3 10.255.2.227/32
300432 P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
300288 P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
300160 P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
300480 P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
300496 P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7

```

```
user@p6> show ldp database
```

```
Input label database, 1.1.1.6:0--1.1.1.2:0
```

```

Label Prefix
3 1.1.1.2/32
299776 1.1.1.3/32
299808 1.1.1.6/32

```

```
Output label database, 1.1.1.6:0--1.1.1.2:0
```

```

Label Prefix
299776 1.1.1.2/32
299792 1.1.1.3/32
3 1.1.1.6/32

```

```
user@pr3> show ldp database
```

Input label database, 1.1.1.3:0--1.1.1.2:0

| Label  | Prefix          |
|--------|-----------------|
| 3      | 1.1.1.2/32      |
| 299776 | 1.1.1.3/32      |
| 299808 | 1.1.1.6/32      |
| 299792 | 10.255.2.227/32 |

Output label database, 1.1.1.3:0--1.1.1.2:0

| Label         | Prefix                                                             |
|---------------|--------------------------------------------------------------------|
| 300096        | 1.1.1.2/32                                                         |
| 3             | 1.1.1.3/32                                                         |
| 299856        | 1.1.1.6/32                                                         |
| 299776        | 10.255.2.227/32                                                    |
| <b>300144</b> | <b>P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7</b>        |
| <b>300128</b> | <b>P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11</b> |

Input label database, 1.1.1.3:0--10.255.2.227:0

| Label  | Prefix          |
|--------|-----------------|
| 300144 | 1.1.1.2/32      |
| 299776 | 1.1.1.3/32      |
| 299856 | 1.1.1.6/32      |
| 3      | 10.255.2.227/32 |

Output label database, 1.1.1.3:0--10.255.2.227:0

| Label  | Prefix          |
|--------|-----------------|
| 300096 | 1.1.1.2/32      |
| 3      | 1.1.1.3/32      |
| 299856 | 1.1.1.6/32      |
| 299776 | 10.255.2.227/32 |

### ***Looking Up the Route Information for the MPLS Label***

**Purpose** Display the point-to-multipoint FEC information.



**Action** user@EgressPE> show route label 299808 detail

```
mpls.0: 14 destinations, 14 routes (14 active, 0 holddown, 0 hidden)
299808 (1 entry, 1 announced)
 *LDP Preference: 9
 Next hop type: Flood
 Address: 0x931922c
 Next-hop reference count: 3
 Next hop type: Router, Next hop index: 1109
 Address: 0x9318b0c
 Next-hop reference count: 2
 Next hop: via so-0/1/3.0
 Label operation: Pop
 Next hop type: Router, Next hop index: 1110
 Address: 0x93191e0
 Next-hop reference count: 2
 Next hop: 192.168.209.11 via fe-1/3/0.0
 Label operation: Pop
 State: **Active Int AckRequest>
 Local AS: 10
 Age: 13:08:15 Metric: 1
 Validation State: unverified
 Task: LDP
 Announcement bits (1): 0-KRT
 AS path: I
 FECs bound to route: P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src:
192.168.219.11
```

### Checking the LDP Traffic Statistics

**Purpose** Monitor the data traffic statistics for the point-to-multipoint LSP.

**Action** user@EgressPE> show ldp traffic-statistics p2mp  
P2MP FEC Statistics:

| FEC(root_addr:lsp_id/grp,src)    | Nexthop     | Packets | Bytes |
|----------------------------------|-------------|---------|-------|
| Shared                           |             |         |       |
| 1.1.1.2:232.2.2.2,1.2.7.7        | so-0/1/3.0  | 0       | 0     |
| No                               |             |         |       |
| 1.1.1.2:232.1.1.1,192.168.219.11 | so-0/1/3.0  | 0       | 0     |
| No                               |             |         |       |
|                                  | fe-1/3/0.0  | 0       | 0     |
| No                               |             |         |       |
| 1.1.1.2:232.1.1.2,192.168.219.11 | so-0/1/3.0  | 0       | 0     |
| No                               |             |         |       |
|                                  | fe-1/3/0.0  | 0       | 0     |
| No                               |             |         |       |
|                                  | lt-1/2/0.14 | 0       | 0     |
| No                               |             |         |       |
| 1.1.1.2:232.1.1.3,192.168.219.11 | fe-1/3/0.0  | 0       | 0     |
| No                               |             |         |       |
| 1.1.1.2:ff3e::1:2,abcd::1:2:7:7  | fe-1/3/0.0  | 0       | 0     |
| No                               |             |         |       |

**Related Documentation**

- *Example: Configuring RSVP-Signaled Point-to-Multipoint LSPs on Logical Systems*
- *Example: Configuring Point-to-Multipoint LDPLSPs as the Data Plane for Intra-AS MBGP MVPNs*

## Configuring Miscellaneous LDP Properties

---

The following sections describe how to configure a number of miscellaneous LDP properties:

- [Configuring LDP to Use the IGP Route Metric on page 628](#)
- [Preventing Addition of Ingress Routes to the inet.0 Routing Table on page 628](#)
- [Multiple-Instance LDP and Carrier-of-Carriers VPNs on page 629](#)
- [Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router on page 629](#)
- [Enabling LDP over RSVP-Established LSPs on page 629](#)
- [Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks on page 630](#)
- [Configuring the TCP MD5 Signature for LDP Sessions on page 630](#)
- [Configuring LDP Session Protection on page 631](#)
- [Disabling SNMP Traps for LDP on page 632](#)
- [Configuring LDP Synchronization with the IGP on LDP Links on page 632](#)
- [Configuring LDP Synchronization with the IGP on the Router on page 633](#)
- [Configuring the Label Withdrawal Timer on page 633](#)
- [Ignoring the LDP Subnet Check on page 633](#)

### Configuring LDP to Use the IGP Route Metric

Use the **track-igp-metric** statement if you want the interior gateway protocol (IGP) route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).

To use the IGP route metric, include the **track-igp-metric** statement:

```
track-igp-metric;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

### Preventing Addition of Ingress Routes to the inet.0 Routing Table

By configuring the **no-forwarding** statement, you can prevent ingress routes from being added to the inet.0 routing table instead of the inet.3 routing table even if you enabled the **traffic-engineering bgp-igp** statement at the **[edit protocols mpls]** or the **[edit logical-systems *logical-system-name* protocols mpls]** hierarchy level. By default, the **no-forwarding** statement is disabled.

To omit ingress routes from the inet.0 routing table, include the **no-forwarding** statement:

```
no-forwarding;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Multiple-Instance LDP and Carrier-of-Carriers VPNs

By configuring multiple LDP routing instances, you can use LDP to advertise labels in a carrier-of-carriers VPN from a service provider provider edge (PE) router to a customer carrier customer edge (CE) router. This is especially useful when the carrier customer is a basic Internet service provider (ISP) and wants to restrict full Internet routes to its PE routers. By using LDP instead of BGP, the carrier customer shields its other internal routers from the Internet. Multiple-instance LDP is also useful when a carrier customer wants to provide Layer 2 or Layer 3 VPN services to its customers.

For an example of how to configure multiple LDP routing instances for carrier-of-carriers VPNs, see the *Multiple Instances for Label Distribution Protocol Feature Guide*.

## Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router

The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop router removes the label and sends the packet to the egress router. If ultimate-hop popping is enabled, label 0 (IPv4 Explicit Null label) is advertised. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label.

To configure ultimate-hop popping, include the **explicit-null** statement:

```
explicit-null;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.



**NOTE:** Juniper Networks routers queue packets based on the incoming label. Routers from other vendors might queue packets differently. Keep this in mind when working with networks containing routers from multiple vendors.

For more information about labels, see “MPLS Label Overview” on page 24 and “MPLS Label Allocation” on page 26.

## Enabling LDP over RSVP-Established LSPs

You can run LDP over LSPs established by RSVP, effectively tunneling the LDP-established LSP through the one established by RSVP. To do so, enable LDP on the lo0.0 interface (see “Enabling and Disabling LDP” on page 528). You must also configure the LSPs over which you want LDP to operate by including the **ldp-tunneling** statement at the [edit protocols mpls label-switched-path *lsp-name*] hierarchy level:

```
[edit]
protocols {
 mpls {
 label-switched-path lsp-name {
 from source;
 to destination;
 ldp-tunneling;
 }
 }
}
```

```
}
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

**Related  
Documentation**

- [Tunneling LDP LSPs in RSVP LSPs Overview on page 523](#)

## Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks

Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This might require that you manually configure the RSVP metric when deploying LDP tunneling over RSVP LSPs in heterogeneous networks.

When a Juniper Networks router is linked to another vendor's router through an RSVP tunnel, and LDP tunneling is also enabled, by default the Juniper Networks router might not use the RSVP tunnel to route traffic to the LDP destinations downstream of the other vendor's egress router if the RSVP path has a metric of 1 larger than the physical OSPF path.

To ensure that LDP tunneling functions properly in heterogeneous networks, you can configure OSPF to ignore the RSVP LSP metric by including the **ignore-lsp-metrics** statement:

```
ignore-lsp-metrics;
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols ospf traffic-engineering shortcuts]**
- **[edit logical-systems *logical-system-name* protocols ospf traffic-engineering shortcuts]**

To enable LDP over RSVP LSPs, you also still need to complete the procedure in Section [“Enabling LDP over RSVP-Established LSPs” on page 629](#).

## Configuring the TCP MD5 Signature for LDP Sessions

You can configure an MD5 signature for an LDP TCP connection to protect against the introduction of spoofed TCP segments into LDP session connection streams.

A router using the MD5 signature option is configured with a password for each peer for which authentication is required. The password is stored encrypted.

LDP hello adjacencies can still be created even when peering interfaces are configured with different security signatures. However, the TCP session cannot be authenticated and is never established.

To configure an MD5 signature for an LDP TCP connection, include the **session** and **authentication-key** statement:

```
session address {
 authentication-key md5-authentication-key;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary section for the **session** statement.

Use the **session** statement to configure the address for the remote end of the LDP session.

The **md5-authentication-key** (password) can be up to 69 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks.

You can also configure an authentication key update mechanism for the LDP routing protocol. This mechanism allows you to update authentication keys without interrupting associated routing and signaling protocols such as Open Shortest Path First (OSPF) and Resource Reservation Setup Protocol (RSVP).

To configure the authentication key update mechanism, include the **key-chain** statement at the **[edit security authentication-key-chains]** hierarchy level, and specify the **key** option to create a keychain consisting of several authentication keys.

```
[edit security authentication-key-chains]
key-chain key-chain-name {
 key key {
 secret secret-data;
 start-time yyyy-mm-dd.hh:mm:ss;
 }
}
```

To configure the authentication key update mechanism for the LDP routing protocol, include the **authentication-key-chain** statement at the **[edit protocols ldp]** hierarchy level to associate the protocol with the **[edit security authentication-key-chains]** authentication keys.

```
[edit protocols ldp]
group group-name {
 neighbor address {
 authentication-key-chain key-chain-name;
 }
}
```

For more information about the authentication key update feature, see *Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols*.

## Configuring LDP Session Protection

An LDP session is normally created between a pair of routers that are connected by one or more links. The routers form one hello adjacency for every link that connects them and associate all the adjacencies with the corresponding LDP session. When the last hello adjacency for an LDP session goes away, the LDP session is terminated. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished.

You can configure the Junos OS to leave the LDP session between two routers up even if there are no hello adjacencies on the links connecting the two routers by configuring the **session-protection** statement. You can optionally specify a time in seconds using the **timeout** option. The session remains up for the duration specified as long as the routers maintain IP network connectivity.

```
session-protection {
```

```
 timeout seconds;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section.

## Disabling SNMP Traps for LDP

Whenever an LDP LSP makes a transition from up to down, or down to up, the router sends an SNMP trap. However, it is possible to disable the LDP SNMP traps on a router, logical system, or routing instance.

For information about the LDP SNMP traps and the proprietary LDP MIB, see the *SNMP MIBs and Traps Reference* and *Interpreting the Enterprise-Specific LDP MIB*.

To disable SNMP traps for LDP, specify the **trap disable** option for the **log-updown** statement:

```
log-updown {
 trap disable;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Configuring LDP Synchronization with the IGP on LDP Links

LDP is a protocol for distributing labels in non-traffic-engineered applications. Labels are distributed along the best path determined by the IGP. If synchronization between LDP and the IGP is not maintained, the LSP goes down. When LDP is not fully operational on a given link (a session is not established and labels are not exchanged), the IGP advertises the link with the maximum cost metric. The link is not preferred but remains in the network topology.

LDP synchronization is supported only on active point-to-point interfaces and LAN interfaces configured as point-to-point under the IGP. LDP synchronization is not supported during graceful restart.

To advertise the maximum cost metric until LDP is operational for synchronization, include the **ldp-synchronization** statement:

```
ldp-synchronization {
 disable;
 hold-time seconds;
}
```

To disable synchronization, include the **disable** statement. To configure the time period to advertise the maximum cost metric for a link that is not fully operational, include the **hold-time** statement.

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

## Configuring LDP Synchronization with the IGP on the Router

You can configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.

To configure the time the LDP waits before informing the IGP that the LDP neighbor and session are operational, include the **igp-synchronization** statement and specify a time in seconds for the **holddown-interval** option:

```
igp-synchronization holddown-interval seconds;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

## Configuring the Label Withdrawal Timer

The label withdrawal timer delays sending a label withdrawal message for a FEC to a neighbor. When an IGP link to a neighbor fails, the label associated with the FEC has to be withdrawn from all the upstream routers if the neighbor is the next hop for the FEC. After the IGP converges and a label is received from a new next hop, the label is readvertised to all the upstream routers. This is the typical network behavior. By delaying label withdrawal by a small amount of time (for example, until the IGP converges and the router receives a new label for the FEC from the downstream next hop), the label withdrawal and sending a label mapping soon could be avoided. The **label-withdrawal-delay** statement allows you to configure this delay time. By default, the delay is 60 seconds.

If the router receives the new label before the timer runs out, the label withdrawal timer is canceled. However, if the timer runs out, the label for the FEC is withdrawn from all of the upstream routers.

By default, LDP waits for 60 seconds before withdrawing labels to avoid resignaling LSPs multiple times while the IGP is reconverging. To configure the label withdrawal delay time in seconds, include the **label-withdrawal-delay** statement:

```
label-withdrawal-delay seconds;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

## Ignoring the LDP Subnet Check

In Junos OS Release 8.4 and later releases, an LDP source address subnet check is performed during the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address. This causes an interoperability issue with some other vendors' equipment.

To disable the subnet check, include the **allow-subnet-mismatch** statement:

```
allow-subnet-mismatch;
```

This statement can be included at the following hierarchy levels:

- [edit protocols ldp **interface** *interface-name*]
- [edit logical-systems *logical-system-name* protocols ldp **interface** *interface-name*]

---

## Configuring LDP LSP Traceroute

You can trace the route followed by an LDP-signaled LSP. LDP LSP traceroute is based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. This feature allows you to periodically trace all paths in a FEC. The FEC topology information is stored in a database accessible from the CLI.

A topology change does not automatically trigger a trace of an LDP LSP. However, you can manually initiate a traceroute. If the traceroute request is for an FEC that is currently in the database, the contents of the database are updated with the results.

The periodic traceroute feature applies to all FECs specified by the **oam** statement configured at the [edit protocols ldp] hierarchy level. To configure periodic LDP LSP traceroute, include the **periodic-traceroute** statement:

```
periodic-traceroute {
 disable;
 exp exp-value;
 fanout fanout-value;
 frequency minutes;
 paths number-of-paths;
 retries retry-attempts;
 source address;
 ttl ttl-value;
 wait seconds;
}
```

You can configure this statement at the following hierarchy levels:

- [edit protocols ldp **oam**]
- [edit protocols ldp **oam** *fec address*]

You can configure the **periodic-traceroute** statement by itself or with any of the following options:

- **exp**—Specify the class of service to use when sending probes.
- **fanout**—Specify the maximum number of next hops to search per node.
- **frequency**—Specify the interval between traceroute attempts.
- **paths**—Specify the maximum number of paths to search.
- **retries**—Specify the number of attempts to send a probe to a specific node before giving up.
- **source**—Specify the IPv4 source address to use when sending probes.



- **ttl**—Specify the maximum time-to-live value. Nodes that are beyond this value are not traced.
- **wait**—Specify the wait interval before resending a probe packet.

## Collecting LDP Statistics

LDP traffic statistics show the volume of traffic that has passed through a particular FEC on a router.

When you configure the **traffic-statistics** statement at the **[edit protocols ldp]** hierarchy level, the LDP traffic statistics are gathered periodically and written to a file. You can configure how often statistics are collected (in seconds) by using the **interval** option. The default collection interval is 5 minutes. You must configure an LDP statistics file; otherwise, LDP traffic statistics are not gathered. If the LSP goes down, the LDP statistics are reset.

To collect LDP traffic statistics, include the **traffic-statistics** statement:

```
traffic-statistics {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 interval interval;
 no-penultimate-hop;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

This section includes the following topics:

- [LDP Statistics Output on page 635](#)
- [Disabling LDP Statistics on the Penultimate-Hop Router on page 636](#)
- [LDP Statistics Limitations on page 637](#)

## LDP Statistics Output

The following sample output is from an LDP statistics file:

| FEC               | Type    | Packets | Bytes | Shared |
|-------------------|---------|---------|-------|--------|
| 10.255.350.448/32 | Transit | 0       | 0     | No     |
|                   | Ingress | 0       | 0     | No     |
| 10.255.350.450/32 | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |
| 10.255.350.451/32 | Transit | 0       | 0     | No     |
|                   | Ingress | 0       | 0     | No     |
| 220.220.220.1/32  | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |
| 220.220.220.2/32  | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |
| 220.220.220.3/32  | Transit | 0       | 0     | Yes    |
|                   | Ingress | 0       | 0     | No     |

May 28 15:02:05, read 12 statistics in 00:00:00 seconds

The LDP statistics file includes the following columns of data:

- **read**—Number of bytes of data passed by the FEC since its LSP came up.
- **read**—FEC for which LDP traffic statistics are collected.
- **read**—Number of packets passed by the FEC since its LSP came up.
- **read**—This number (which appears next to the date and time) might differ from the actual number of the statistics displayed. Some of the statistics are summarized before being displayed.
- **Shared**—A **Yes** value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such.
- **Type**—Type of traffic originating from a router, either **Ingress** (originating from this router) or **Transit** (forwarded through this router).

## Disabling LDP Statistics on the Penultimate-Hop Router

Gathering LDP traffic statistics at the penultimate-hop router can consume excessive system resources, on next-hop routes in particular. This problem is exacerbated if you have configured the **deaggregate** statement in addition to the **traffic-statistics** statement. For routers reaching their limit of next-hop route usage, we recommend configuring the **no-penultimate-hop** option for the **traffic-statistics** statement:

```
traffic-statistics {
 no-penultimate-hop;
}
```

For a list of hierarchy levels at which you can configure the **traffic-statistics** statement, see the statement summary section for this statement.



**NOTE:** When you configure the **no-penultimate-hop** option, no statistics are available for the FECs that are the penultimate hop for this router.

Whenever you include or remove this option from the configuration, the LDP sessions are taken down and then restarted.

The following sample output is from an LDP statistics file showing routers on which the **no-penultimate-hop** option is configured:

| FEC               | Type    | Packets             | Bytes | Shared |
|-------------------|---------|---------------------|-------|--------|
| 10.255.245.218/32 | Transit | 0                   | 0     | No     |
|                   | Ingress | 4                   | 246   | No     |
| 10.255.245.221/32 | Transit | statistics disabled |       |        |
|                   | Ingress | statistics disabled |       |        |
| 13.1.1.0/24       | Transit | statistics disabled |       |        |
|                   | Ingress | statistics disabled |       |        |
| 13.1.3.0/24       | Transit | statistics disabled |       |        |
|                   | Ingress | statistics disabled |       |        |

## LDP Statistics Limitations

The following are issues related to collecting LDP statistics by configuring the **traffic-statistics** statement:

- You cannot clear the LDP statistics.
- If you shorten the specified interval, a new LDP statistics request is issued only if the statistics timer expires later than the new interval.
- A new LDP statistics collection operation cannot start until the previous one has finished. If the interval is short or if the number of LDP statistics is large, the time gap between the two statistics collections might be longer than the interval.

When an LSP goes down, the LDP statistics are reset.

## Tracing LDP Protocol Traffic

The following sections describe how to configure the trace options to examine LDP protocol traffic:

- [Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels on page 637](#)
- [Tracing LDP Protocol Traffic Within FECs on page 638](#)
- [Examples: Tracing LDP Protocol Traffic on page 639](#)

### Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels

To trace LDP protocol traffic, you can specify options in the global **traceoptions** statement at the **[edit routing-options]** hierarchy level, and you can specify LDP-specific options by including the **traceoptions** statement:

```
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

Use the **file** statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory `/var/log`. We recommend that you place LDP-tracing output in the file **ldp-log**.

The following trace flags display the operations associated with the sending and receiving of various LDP messages. Each can carry one or more of the following modifiers:

- **address**—Trace the operation of address and address withdrawal messages.
- **binding**—Trace label-binding operations.
- **error**—Trace error conditions.
- **event**—Trace protocol events.

- **initialization**—Trace the operation of initialization messages.
- **label**—Trace the operation of label request, label map, label withdrawal, and label release messages.
- **notification**—Trace the operation of notification messages.
- **packets**—Trace the operation of address, address withdrawal, initialization, label request, label map, label withdrawal, label release, notification, and periodic messages. This modifier is equivalent to setting the **address**, **initialization**, **label**, **notification**, and **periodic** modifiers.

You can also configure the **filter** flag modifier with the **match-on address** sub-option for the **packets** flag. This allows you to trace based on the source and destination addresses of the packets.

- **path**—Trace label-switched path operations.
- **path**—Trace label-switched path operations.
- **periodic**—Trace the operation of hello and keepalive messages.
- **route**—Trace the operation of route messages.
- **state**—Trace protocol state transitions.

## Tracing LDP Protocol Traffic Within FECs

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers.

You can trace LDP protocol traffic within a specific FEC and filter LDP trace statements based on an FEC. This is useful when you want to trace or troubleshoot LDP protocol traffic associated with an FEC. The following trace flags are available for this purpose: **route**, **path**, and **binding**.

The following example illustrates how you might configure the LDP **traceoptions** statement to filter LDP trace statements based on an FEC:

```
[edit protocols ldp traceoptions]
set flag route filter match-on fec policy "filter-policy-for-ldp-fec";
```

This feature has the following limitations:

- The filtering capability is only available for FECs composed of IP version 4 (IPv4) prefixes.
- Layer 2 circuit FECs cannot be filtered.
- When you configure both route tracing and filtering, MPLS routes are not displayed (they are blocked by the filter).

- Filtering is determined by the policy and the configured value for the **match-on** option. When configuring the policy, be sure that the default behavior is always **reject**.
- The only **match-on** option is **fec**. Consequently, the only type of policy you should include is a route-filter policy.

## Examples: Tracing LDP Protocol Traffic

Trace LDP path messages in detail:

```
[edit]
protocols {
 ldp {
 traceoptions {
 file ldp size 10m files 5;
 flag path;
 }
 }
}
```

Trace all LDP outgoing messages:

```
[edit]
protocols {
 ldp {
 traceoptions {
 file ldp size 10m files 5;
 flag packets;
 }
 }
}
```

Trace all LDP error conditions:

```
[edit]
protocols {
 ldp {
 traceoptions {
 file ldp size 10m files 5;
 flag error;
 }
 }
}
```

Trace all LDP incoming messages and all label-binding operations:

```
[edit]
protocols {
 ldp {
 traceoptions {
 file ldp size 10m files 5 world-readable;
 flag packets receive;
 flag binding;
 }
 interface all {
 }
 }
}
```

```
}
```

Trace LDP protocol traffic for an FEC associated with the LSP:

```
[edit]
protocols {
 ldp {
 traceoptions {
 flag route filter match-on fec policy filter-policy-for-ldp-fec;
 }
 }
}
```

## PART 5

# Configuring Circuit Cross-Connect (CCC) and Translational Cross-Connect (TCC)

- [CCC and TCC Overview on page 643](#)
- [Configuring CCC and TCC on page 647](#)





## CHAPTER 15

# CCC and TCC Overview

- [CCC Overview on page 643](#)
- [Transmitting Nonstandard BPDUs on page 644](#)
- [TCC Overview on page 644](#)

### CCC Overview

---

Circuit cross-connect (CCC) allows you to configure transparent connections between two circuits, where a circuit can be a Frame Relay data-link connection identifier (DLCI), an Asynchronous Transfer Mode (ATM) virtual circuit (VC), a Point-to-Point Protocol (PPP) interface, a Cisco High-Level Data Link Control (HDLC) interface, or an MPLS label-switched path (LSP). Using CCC, packets from the source circuit are delivered to the destination circuit with, at most, the Layer 2 address being changed. No other processing—such as header checksums, time-to-live (TTL) decrementing, or protocol processing—is done.

CCC circuits fall into two categories: logical interfaces, which include DLCIs, VCs, virtual local area network (VLAN) IDs, PPP and Cisco HDLC interfaces, and LSPs. The two circuit categories provide three types of cross-connect:

- **Layer 2 switching**—Cross-connects between logical interfaces provide what is essentially Layer 2 switching. The interfaces that you connect must be of the same type.
- **MPLS tunneling**—Cross-connects between interfaces and LSPs allow you to connect two distant interface circuits of the same type by creating MPLS tunnels that use LSPs as the conduit.
- **LSP stitching**—Cross-connects between LSPs provide a way to “stitch” together two label-switched paths, including paths that fall in two different traffic engineering database areas.

For Layer 2 switching and MPLS tunneling, the cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first. For LSP stitching, the cross-connect is unidirectional.

You can police (control) the amount of traffic flowing over CCC circuits. For more information, see the *Junos OS VPNs Library for Routing Devices*.

It is also possible to use the **ping** command to check the integrity of CCC LSPs. See [“Pinging CCC LSPs” on page 360](#) for more information.

## Transmitting Nonstandard BPDUs

CCC protocol (and Layer 2 Circuit and Layer 2 VPN) configurations can transmit nonstandard bridge protocol data units (BPDUs) generated by other vendors' equipment. This is the default behavior on all supported PICs and requires no additional configuration.


The following PICs are supported on M320 and T Series routers:

- 1-port Gigabit Ethernet PIC
- 2-port Gigabit Ethernet PIC
- 4-port Gigabit Ethernet PIC
- 10-port Gigabit Ethernet PIC

## TCC Overview

Translational cross-connect (TCC) is a switching concept that enables you to establish interconnections between a variety of Layer 2 protocols or circuits. It is similar to CCC. However, whereas CCC requires the same Layer 2 encapsulations on each side of a Juniper Networks router (such as PPP-to-PPP or Frame Relay-to-Frame Relay), TCC enables you to connect different types of Layer 2 protocols interchangeably. When you use TCC, combinations such as PPP-to-ATM (see [Figure 65 on page 644](#)) and Ethernet-to-Frame Relay connections are possible.

Figure 65: TCC Example



The Layer 2 circuits and encapsulation types that can be interconnected by TCC are:

- Ethernet
- Extended VLANs
- PPP
- HDLC
- ATM
- Frame Relay

TCC works by removing the Layer 2 header when frames enter the router and adding a different Layer 2 header on the frames before they leave the router. In [Figure 65 on page 644](#), the PPP encapsulation is stripped from the frames arriving at Router B, and the ATM encapsulation is added before the frames are sent to Router C.

Note that all control traffic is terminated at the interconnecting router (Router B). Examples of traffic controllers include the Link Control Protocol (LCP) and the Network Control Protocol (NCP) for PPP, keepalives for HDLC, and Local Management Interface (LMI) for Frame Relay.

TCC functionality is different from standard Layer 2 switching. TCC only swaps Layer 2 headers. No other processing, such as header checksums, TTL decrementing, or protocol handling is performed. TCC is supported for IPv4 only.

Address Resolution Protocol (APR) packet policing on TCC Ethernet interfaces is effective for releases 10.4 and onwards.

You can configure TCC for interface switching and for Layer 2 VPNs. For more information about using TCC for virtual private networks (VPNs), see the *Junos OS VPNs Library for Routing Devices*.



# Configuring CCC and TCC

- [Configuring Layer 2 Switching Cross-Connects Using CCC on page 647](#)
- [Configuring MPLS LSP Tunnel Cross-Connects Using CCC on page 655](#)
- [Configuring TCC on page 659](#)
- [CCC and TCC Graceful Restart on page 664](#)
- [Configuring CCC and TCC Graceful Restart on page 665](#)
- [Configuring CCC Switching for Point-to-Multipoint LSPs on page 665](#)


## Configuring Layer 2 Switching Cross-Connects Using CCC

Layer 2 switching cross-connects join logical interfaces to form what is essentially Layer 2 switching. The interfaces that you connect must be of the same type.

[Figure 66 on page 647](#) illustrates a Layer 2 switching cross-connect. In this topology, Router A and Router C have Frame Relay connections to Router B, which is a Juniper Networks router. Circuit cross-connect (CCC) allows you to configure Router B to act as a Frame Relay (Layer 2) switch.

To configure Router B to act as a Frame Relay switch, you configure a circuit from Router A to Router C that passes through Router B, effectively configuring Router B as a Frame Relay switch with respect to these routers. This configuration allows Router B to transparently switch packets (frames) between Router A and Router C without regard to the packets' contents or the Layer 3 protocols. The only processing that Router B performs is to translate DLCI 600 to 750.

**Figure 66: Layer 2 Switching Cross-Connect**



If the Router A-to-Router B and Router B-to-Router C circuits were PPP, for example, the Link Control Protocol and Network Control Protocol exchanges occur between Router A and Router C. These messages are handled transparently by Router B, allowing Router A and Router C to use various PPP options (such as header or address compression and authentication) that Router B might not support. Similarly, Router A and Router C exchange keepalives, providing circuit-to-circuit connectivity status.

You can configure Layer 2 switching cross-connects on PPP, Cisco HDLC, Frame Relay, Ethernet, and ATM circuits. In a single cross-connect, only like interfaces can be connected.

To configure Layer 2 switching cross-connects, you must configure the following on the router that is acting as the switch (Router B in [Figure 66 on page 647](#)):

- [Configuring the CCC Encapsulation for Layer 2 Switching Cross-Connects on page 648](#)
- [Configuring the CCC Connection for Layer 2 Switching Cross-Connects on page 652](#)
- [Configuring MPLS for Layer 2 Switching Cross-Connects on page 652](#)
- [Example: Configuring a Layer 2 Switching Cross-Connect on page 653](#)

## Configuring the CCC Encapsulation for Layer 2 Switching Cross-Connects

To configure Layer 2 switching cross-connects, configure the CCC encapsulation on the router that is acting as the switch (Router B in [Figure 66 on page 647](#)).



**NOTE:** You cannot configure families on CCC interfaces; that is, you cannot include the family statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level.

For instructions for configuring the encapsulation for Layer 2 switching cross-connects, see the following sections:

- [Configuring ATM Encapsulation for Layer 2 Switching Cross-Connects on page 648](#)
- [Configuring Ethernet Encapsulation for Layer 2 Switching Cross-Connects on page 649](#)
- [Configuring Ethernet VLAN Encapsulation for Layer 2 Switching Cross-Connects on page 649](#)
- [Configuring Aggregated Ethernet Encapsulation for Layer 2 Switching Cross-Connects on page 650](#)
- [Configuring Frame Relay Encapsulation for Layer 2 Switching Cross-Connects on page 651](#)
- [Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching Cross-Connects on page 652](#)

## Configuring ATM Encapsulation for Layer 2 Switching Cross-Connects

---

For ATM circuits, specify the encapsulation when configuring the virtual circuit (VC). Configure each VC as a circuit or a regular logical interface by including the following statements:

```
at-fpc/pic/port {
 atm-options {
 vpi vpi-identifier maximum-vcs maximum-vcs;
 }
 unit logical-unit-number {
 encapsulation encapsulation-type;
 point-to-point; # Default interface type
 vci vpi-identifier.vci-identifier;
```

```
}
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

### Configuring Ethernet Encapsulation for Layer 2 Switching Cross-Connects

For Ethernet circuits, specify **ethernet-ccc** in the **encapsulation** statement. This statement configures the entire physical device. For these circuits to work, you must also configure a logical interface (unit 0).

Ethernet interfaces with standard Tag Protocol Identifier (TPID) tagging can use Ethernet CCC encapsulation. On M Series Multiservice Edge Routers, except the M320, one-port Gigabit Ethernet, two-port Gigabit Ethernet, four-port Gigabit Ethernet, and four-port Fast Ethernet PICs can use Ethernet CCC encapsulation. On T Series Core Routers and M320 routers, one-port Gigabit Ethernet and two-port Gigabit Ethernet PICs installed in FPC2 can use Ethernet CCC encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.

```
fe-fpc/pic/port {
 encapsulation ethernet-ccc;
 unit 0;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

### Configuring Ethernet VLAN Encapsulation for Layer 2 Switching Cross-Connects

An Ethernet virtual LAN (VLAN) circuit can be configured using either the **vlan-ccc** or **extended-vlan-ccc** encapsulation. If you configure the **extended-vlan-ccc** encapsulation on the physical interface, you cannot configure the **inet** family on the logical interfaces. Only the **ccc** family is allowed. If you configure the **vlan-ccc** encapsulation on the physical interface, both the **inet** and **ccc** families are supported on the logical interfaces. Ethernet interfaces in VLAN mode can have multiple logical interfaces.

For encapsulation type **vlan-ccc**, VLAN IDs from 512 through 4094 are reserved for CCC VLANs. For the **extended-vlan-ccc** encapsulation type, all VLAN IDs 1 and higher are valid. VLAN ID 0 is reserved for tagging the priority of frames.



**NOTE:** Some vendors use the proprietary TPIDs 0x9100 and 0x9901 to encapsulate a VLAN-tagged packet into a VLAN-CCC tunnel to interconnect a geographically separated metro Ethernet network. By configuring the **extended-vlan-ccc** encapsulation type, a Juniper Networks router can accept all three TPIDs (0x8100, 0x9100, and 0x9901).

Configure an Ethernet VLAN circuit with the **vlan-ccc** encapsulation as follows:

```
interfaces {
 type-fpc/pic/port {
 vlan-tagging;
 encapsulation vlan-ccc;
 unit logical-unit-number {
 encapsulation vlan-ccc;
 vlan-id vlan-id;
 }
 }
}
```

You can configure these statements at the following hierarchy levels:

- **[edit interfaces]**
- **[edit logical-systems *logical-system-name* interfaces]**

Configure an Ethernet VLAN circuit with the **extended-vlan-ccc** encapsulation statement as follows:

```
interfaces {
 type-fpc/pic/port {
 vlan-tagging;
 encapsulation extended-vlan-ccc;
 unit logical-unit-number {
 vlan-id vlan-id;
 family ccc;
 }
 }
}
```

You can configure these statements at the following hierarchy levels:

- **[edit interfaces]**
- **[edit logical-systems *logical-system-name* interfaces]**

Whether you configure the encapsulation as **vlan-ccc** or **extended-vlan-ccc**, you must enable VLAN tagging by including the **vlan-tagging** statement.

### Configuring Aggregated Ethernet Encapsulation for Layer 2 Switching Cross-Connects

---

You can configure aggregated Ethernet interfaces for CCC connections and for Layer 2 virtual private networks (VPNs).

Aggregated Ethernet interfaces configured with VLAN tagging can be configured with multiple logical interfaces. The only encapsulation available for aggregated Ethernet logical interfaces is **vlan-ccc**. When you configure the **vlan-id** statement, you are limited to VLAN IDs 512 through 4094.

Aggregated Ethernet interfaces configured without VLAN tagging can be configured only with the **ethernet-ccc** encapsulation. All untagged Ethernet packets received are forwarded based on the CCC parameters.



To configure aggregated Ethernet interfaces for CCC connections, include the **ae0** statement at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
ae0 {
 encapsulation (ethernet-ccc | extended-vlan-ccc | vlan-ccc);
 vlan-tagging;
 aggregated-ether-options {
 minimum-links links;
 link-speed speed;
 }
 unit logical-unit-number {
 encapsulation vlan-ccc;
 vlan-id identifier;
 family ccc;
 }
}
```

Be aware of the following limitations when configuring CCC connections over aggregated Ethernet interfaces:

- If you configured load balancing between child links, be aware that a different hash key is used to distribute packets among the child links. Standard aggregated interfaces have family inet configured. An IP version 4 (IPv4) hash key (based on the Layer 3 information) is used to distribute packets among the child links. A CCC connection over an aggregated Ethernet interface has family ccc configured instead. Instead of an IPv4 hash key, an MPLS hash key (based on the destination media access control [MAC] address) is used to distributed packets among the child links.
- The extended-vlan-ccc encapsulation is not supported on the 12-port Fast Ethernet PIC and the 48-port Fast Ethernet PIC.
- The Junos OS does not support the Link Aggregation Control Protocol (LACP) when an aggregated interface is configured as a VLAN (with vlan-ccc encapsulation). LACP can be configured only when the aggregated interface is configured with the ethernet-ccc encapsulation.

For more information about how to configure aggregated Ethernet interfaces, see the *Junos OS Network Interfaces Library for Routing Devices*.

### Configuring Frame Relay Encapsulation for Layer 2 Switching Cross-Connects

For Frame Relay circuits, specify the encapsulation when configuring the DLCI. Configure each DLCI as a circuit or a regular logical interface. The DLCI for regular interfaces must be from 1 through 511. For CCC interfaces, it must be from 512 through 4094.

```
interfaces {
 type-fpc/pic/port {
 unit logical-unit-number {
 dlci dlci-identifier;
 encapsulation encapsulation-type;
 point-to-point; # Default interface type
 }
 }
}
```

You can configure these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

### Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching Cross-Connects

---

For PPP and Cisco HDLC circuits, specify the encapsulation in the **encapsulation** statement. This statement configures the entire physical device. For these circuits to work, you must configure a logical interface (unit 0).

```
interfaces type-fpc/pic/port {
 encapsulation encapsulation-type;
 unit 0;
}
```

You can configure these statements at the following hierarchy levels:

- [edit interfaces *type-fpc/pic/port*]
- [edit logical-systems *logical-system-name* interfaces *type-fpc/pic/port*]

### Configuring the CCC Connection for Layer 2 Switching Cross-Connects

To configure Layer 2 switching cross-connects, define the connection between the two circuits by including the **interface-switch** statement. You configure this connection on the router that is acting as the switch (Router B in [Figure 66 on page 647](#)). The connection joins the interface that comes from the circuit's source to the interface that leads to the circuit's destination. When you specify the interface names, include the logical portion of the name, which corresponds to the logical unit number. The cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first.

```
interface-switch connection-name {
 interface interface-name.unit-number;
 interface interface-name.unit-number;
}
```

You can include this statement at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

### Configuring MPLS for Layer 2 Switching Cross-Connects

For Layer 2 switching cross-connects to work, you must enable MPLS on the router by including at least the following statements. This minimum configuration enables MPLS on a logical interface for the switching cross-connect.

Include the **family mpls** statement:

```
family mpls;
```

You can configure this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

You can then specify this logical interface in the MPLS protocol configuration:

```
mpls {
 interface interface-name; # Required to enable MPLS on the interface
}
```


You can configure these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

### Example: Configuring a Layer 2 Switching Cross-Connect

Configure a full-duplex Layer 2 switching cross-connect between Router A and Router C, using a Juniper Networks router, Router B, as the virtual switch. See the topology in [Figure 67 on page 653](#) and [Figure 68 on page 654](#).

Figure 67: Topology of a Frame Relay Layer 2 Switching Cross-Connect




```
[edit]
interfaces {
 so-1/0/0 {
 encapsulation frame-relay-ccc;
 unit 1 {
 point-to-point;
 encapsulation frame-relay-ccc;
 dlci 600;
 }
 }
 so-2/0/0 {
 encapsulation frame-relay-ccc;
 unit 2 {
 point-to-point;
 encapsulation frame-relay-ccc;
 dlci 750;
 }
 }
}
protocols {
 connections {
 interface-switch router-a-to-router-c {
 interface so-1/0/0.1;
 }
 }
}
```

```

 interface so-2/0/0.2;
 }
}
mpls {
 interface all;
}
}

```

Figure 68: Sample Topology of a VLAN Layer 2 Switching Cross-Connect



```


[edit]
interfaces {
 ge-2/1/0 {
 vlan-tagging;
 encapsulation vlan-ccc;
 unit 0 {
 encapsulation vlan-ccc;
 vlan-id 600;
 }
 }
 ge-2/2/0 {
 vlan-tagging;
 encapsulation vlan-ccc;
 unit 0 {
 encapsulation vlan-ccc;
 vlan-id 600;
 }
 unit 1 {
 family inet {
 vlan-id 1;
 address 10.9.200.1/24;
 }
 }
 }
}
protocols {
 mpls {
 interface all;
 }
 connections {
 interface-switch layer2-sw {
 interface ge-2/1/0.0;
 interface ge-2/2/0.0;
 }
 }
}
}

```

## Configuring MPLS LSP Tunnel Cross-Connects Using CCC

MPLS tunnel cross-connects between interfaces and LSPs allow you to connect two distant interface circuits of the same type by creating MPLS tunnels that use LSPs as the conduit. The topology in [Figure 69 on page 655](#) illustrates an MPLS LSP tunnel cross-connect. In this topology, two separate networks, in this case ATM access networks, are connected through an IP backbone. CCC allows you to establish an LSP tunnel between the two domains. With LSP tunneling, you tunnel the ATM traffic from one network across a SONET backbone to the second network by using an MPLS LSP.

**Figure 69: MPLS Tunnel Cross-Connect**



When traffic from Router A (VC 234) reaches Router B, it is encapsulated and placed into an LSP, which is sent through the backbone to Router C. At Router C, the label is removed, and the packets are placed onto the ATM permanent virtual circuit (PVC) (VC 591) and sent to Router D. Similarly, traffic from Router D (VC 591) is sent over an LSP to Router B, then placed on VC 234 to Router A.

You can configure LSP tunnel cross-connect on PPP, Cisco HDLC, Frame Relay, and ATM circuits. In a single cross-connect, only like interfaces can be connected.

When you use MPLS tunnel cross-connects to support IS-IS, you must ensure that the LSP's maximum transmission unit (MTU) can, at a minimum, accommodate a 1492-octet IS-IS protocol data unit (PDU) in addition to the link-level overhead associated with the technology being connected.

For the tunnel cross-connects to work, the IS-IS frame size on the edge routers (Routers A and D in [Figure 70 on page 658](#)) must be smaller than the LSP's MTU.



**NOTE:** Frame size values do not include the frame check sequence (FCS) or delimiting flags.

To determine the LSP MTU required to support IS-IS, use the following calculation:

$$\text{IS-IS MTU (minimum 1492, default 1497) + frame overhead + 4 (MPLS shim header) = Minimum LSP MTU}$$

The framing overhead varies based on the encapsulation being used. The following lists the IS-IS encapsulation overhead values for various encapsulations:

- ATM
  - AAL5 multiplex—8 bytes (RFC 1483)

- VC multiplex—0 bytes
- Frame Relay
  - Multiprotocol—2 bytes (RFCs 1490 and 2427)
  - VC multiplex—0 bytes
- HDLC—4 bytes
- PPP—4 bytes
- VLAN—21 bytes (802.3/LLC)

For IS-IS to work over VLAN-CCC, the LSP's MTU must be at least 1513 bytes (or 1518 for 1497-byte PDUs). If you increase the size of a Fast Ethernet MTU above the default of 1500 bytes, you might need to explicitly configure jumbo frames on intervening equipment.

To modify the MTU, include the **mtu** statement when configuring the logical interface family at the **[edit interfaces *interface-name* unit *logical-unit-number* encapsulation *family*]** hierarchy level. For more information about setting the MTU, see the *Junos OS Network Interfaces Library for Routing Devices*.

To configure an LSP tunnel cross-connect, you must configure the following on the interdomain router (Router B in [Figure 70 on page 658](#)):

- [Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects on page 656](#)
- [Configuring the CCC Connection for LSP Tunnel Cross-Connects on page 657](#)
- [Example: Configuring an LSP Tunnel Cross-Connect on page 658](#)

## Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects

To configure LSP tunnel cross-connects, you must configure the CCC encapsulation on the ingress and egress routers (Router B and Router C, respectively, in [Figure 70 on page 658](#)).



**NOTE:** You cannot configure families on CCC interfaces; that is, you cannot include the **family** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

For PPP or Cisco HDLC circuits, include the **encapsulation** statement to configure the entire physical device. For these circuits to work, you must configure logical unit 0 on the interface.

```
type-fpc/pic/port {
 encapsulation (ppp-ccc | cisco-hdlc-ccc);
 unit 0;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces]**

- [edit logical-systems *logical-system-name* interfaces]

For ATM circuits, specify the encapsulation when configuring the VC by including the following statements. For each VC, you configure whether it is a circuit or a regular logical interface.

```
at-fpc/pic/port {
 atm-options {
 vpi vpi-identifier maximum-vcs maximum-vcs;
 }
 unit logical-unit-number {
 point-to-point; # Default interface type
 encapsulation atm-ccc-vc-mux;
 vci vpi-identifier.vci-identifier;
 }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

For Frame Relay circuits, include the following statements to specify the encapsulation when configuring the DLCI. For each DLCI, you configure whether it is a circuit or a regular logical interface. The DLCI for regular interfaces must be in the range 1 through 511. For CCC interfaces, it must be in the range 512 through 1022.

```
type-fpc/pic/port {
 encapsulation frame-relay-ccc;
 unit logical-unit-number {
 point-to-point; # default interface type
 encapsulation frame-relay-ccc;
 dlci dlci-identifier;
 }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

For more information about the **encapsulation** statement, see the *Junos OS Network Interfaces Library for Routing Devices*.

## Configuring the CCC Connection for LSP Tunnel Cross-Connects

To configure LSP tunnel cross-connects, include the **remote-interface-switch** statement to define the connection between the two circuits on the ingress and egress routers (Router B and Router C, respectively, in [Figure 70 on page 658](#)). The connection joins the interface or LSP that comes from the circuit's source to the interface or LSP that leads to the circuit's destination. When you specify the interface name, include the logical portion of the name, which corresponds to the logical unit number. For the cross-connect to be bidirectional, you must configure cross-connects on two routers.

```
remote-interface-switch connection-name {
 interface interface-name.unit-number;
 transmit-lsp label-switched-path;
 receive-lsp label-switched-path;
}
```


You can include these statements at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

### Example: Configuring an LSP Tunnel Cross-Connect

Configure a full-duplex MPLS LSP tunnel cross-connect from Router A to Router D, passing through Router B and Router C. See the topology in [Figure 70 on page 658](#).

Figure 70: Example Topology of MPLS LSP Tunnel Cross-Connect



On Router B:

```
[edit]
interfaces {
 at-7/1/1 {
 atm-options {
 vpi 1 maximum-vcs 600;
 }
 unit 1 {
 point-to-point; # default interface type
 encapsulation atm-ccc-vc-mux;
 vci 1.234;
 }
 }
}
protocols {
 connections {
 remote-interface-switch router-b-to-router-c {
 interface at-7/1/1.1;
 transmit-lsp lsp1;
 receive-lsp lsp2;
 }
 }
}
```

On Router C:

```
[edit]
interfaces {
 at-3/0/0 {
 atm-options {
```



```

 vpi 2 maximum-vcs 600;
 }
 unit 2 {
 point-to-point; # default interface type
 encapsulation atm-ccc-vc-mux;
 vci 2.591;
 }
}
protocols {
 connections {
 remote-interface-switch router-b-to-router-c {
 interface at-3/0/0.2;
 transmit-lsp lsp2;
 receive-lsp lsp1;
 }
 }
}

```

## Configuring TCC

This section describes how to configure translational cross-connect (TCC). Extensive examples on how to configure TCC for interface switching and for Layer 2.5 VPNs are available in the *Junos OS, Release 15.1*.

To configure TCC, you must perform the following tasks on the router that is acting as the switch:

- [Configuring the Encapsulation for Layer 2 Switching TCCs on page 659](#)
- [Configuring the Connection for Layer 2 Switching TCCs on page 663](#)
- [Configuring MPLS for Layer 2 Switching TCCs on page 663](#)

### Configuring the Encapsulation for Layer 2 Switching TCCs

To configure a Layer 2 switching TCC, specify the TCC encapsulation on the desired interfaces of the router that is acting as the switch.



**NOTE:** You cannot configure standard protocol families on TCC or CCC interfaces. Only the CCC family is allowed on CCC interfaces, and only the TCC family is allowed on TCC interfaces.

For Ethernet circuits and Ethernet extended VLAN circuits, you must also configure the Address Resolution Protocol (ARP). See [“Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations” on page 662](#).

- [Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching TCCs on page 660](#)
- [Configuring ATM Encapsulation for Layer 2 Switching TCCs on page 660](#)
- [Configuring Frame Relay Encapsulation for Layer 2 Switching TCCs on page 660](#)
- [Configuring Ethernet Encapsulation for Layer 2 Switching TCCs on page 661](#)

- [Configuring Ethernet Extended VLAN Encapsulation for Layer 2 Switching TCCs on page 661](#)
- [Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations on page 662](#)

---

### Configuring PPP and Cisco HDLC Encapsulation for Layer 2 Switching TCCs

For PPP and Cisco HDLC circuits, configure the encapsulation type for the entire physical device by specifying the appropriate value for the **encapsulation** statement. For these circuits to work, you must also configure the logical interface **unit 0**.

```
encapsulation (ppp-tcc | cisco-hdlc-tcc);
unit 0{...}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name*]
- [edit logical-systems *logical-system-name* interfaces *interface-name*]

---

### Configuring ATM Encapsulation for Layer 2 Switching TCCs

For ATM circuits, configure the encapsulation type by specifying the appropriate value for the **encapsulation** statement in the virtual circuit (VC) configuration. Specify whether each VC is a circuit or a regular logical interface.

```
atm-options {
 vpi vpi-identifier maximum-vcs maximum-vcs;
}
unit logical-unit-number {
 encapsulation (atm-tcc-vc-mux | atm-tcc-snap);
 point-to-point;
 vci vpi-identifier.vci-identifier;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces at-*fpc/pic/port*]
- [edit logical-systems *logical-system-name* interfaces at-*fpc/pic/port*]

---

### Configuring Frame Relay Encapsulation for Layer 2 Switching TCCs

For Frame Relay circuits, configure the encapsulation type by specifying the value **frame-relay-tcc** for the **encapsulation** statement when configuring the data-link connection identifier (DLCI). You configure each DLCI as a circuit or a regular logical interface. The DLCI for regular interfaces must be in the range from 1 through 511, but for TCC and CCC interfaces it must be in the range from 512 through 1022.

```
encapsulation frame-relay-tcc;
unit logical-unit-number {
 dlci dlci-identifier;
 encapsulation frame-relay-tcc;
 point-to-point;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name*]
- [edit logical-systems *logical-system-name* interfaces *interface-name*]

### Configuring Ethernet Encapsulation for Layer 2 Switching TCCs

For Ethernet TCC circuits, configuring the encapsulation type for the entire physical device by specifying the value **ethernet-tcc** for the **encapsulation** statement.

You must also specify static values for a remote address and a proxy address at the [edit interfaces *interface-name* unit *unit-number* family **tcc**] or [edit logical-systems *logical-system-name* interfaces *interface-name* unit *unit-number* family **tcc**] hierarchy level.

The remote address is associated with the TCC switching router's Ethernet neighbor; in the **remote** statement you must specify both the IP address and the media access control (MAC) address of the Ethernet neighbor. The proxy address is associated with the TCC router's other neighbor connected by the unlike link; in the **proxy** statement you must specify the IP address of the non-Ethernet neighbor.

You can configure Ethernet TCC encapsulation for the interfaces on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Fast Ethernet, and 4-port Gigabit Ethernet PICs.

```
encapsulation ethernet-tcc;
unit logical-unit-number {
 family tcc {
 proxy {
 inet-address ip-address;
 }
 remote {
 inet-address ip-address;
 mac-address mac-address;
 }
 }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces (fe | ge)-*fpc/pic/port*]
- [edit logical-systems *logical-system-name* interfaces (fe | ge)-*fpc/pic/port*]



**NOTE:** For Ethernet circuits, you must also configure the Address Resolution Protocol (ARP). See “[Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations](#)” on page 662.

### Configuring Ethernet Extended VLAN Encapsulation for Layer 2 Switching TCCs

For Ethernet extended VLAN circuits, configure the encapsulation type for the entire physical device by specifying the value **extended-vlan-tcc** for the **encapsulation** statement.

You must also enable VLAN tagging. Ethernet interfaces in VLAN mode can have multiple logical interfaces. With encapsulation type **extended-vlan-tcc**, all VLAN IDs from 0 through

4094 are valid, up to a maximum of 1024 VLANs. As with Ethernet circuits, you must also specify a proxy address and a remote address at the `[edit interfaces interface-name unit logical-unit-number family tcc]` or `[edit logical-systems logical-system-name interfaces interface-name unit unit-number family tcc]` hierarchy level (see “[Configuring Ethernet Encapsulation for Layer 2 Switching TCCs](#)” on page 661).

```
encapsulation extended-vlan-tcc;
vlan-tagging;
unit logical-unit-number {
 vlan-id identifier;
 family tcc;
 proxy {
 inet-address ip-address;
 }
 remote {
 inet-address ip-address;
 mac-address mac-address;
 }
}
```

You can configure these statements at the following hierarchy levels:

- `[edit interfaces interface-name]`
- `[edit logical-systems logical-system-name interfaces interface-name]`



**NOTE:** For Ethernet extended VLAN circuits, you must also configure the Address Resolution Protocol (ARP). See “[Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations](#)” on page 662.

---

### Configuring ARP for Ethernet and Ethernet Extended VLAN Encapsulations

---

For Ethernet and Ethernet extended VLAN circuits with TCC encapsulation, you must also configure ARP. Because TCC simply removes one Layer 2 header and adds another, the default form of dynamic ARP is not supported; you must configure static ARP.

Because remote and proxy addresses are specified on the router performing TCC switching, you must apply the static ARP statement to the Ethernet-type interfaces of the routers that connect to the TCC-switched router. The `arp` statement must specify the IP address and the MAC address of the remotely connected neighbor by use of the unlike Layer 2 protocol on the far side of the TCC switching router.

```
arp ip-address mac mac-address;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number family inet address ip-address]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet address ip-address]`

## Configuring the Connection for Layer 2 Switching TCCs

You must configure the connection between the two circuits of the Layer 2 switching TCC on the router acting as the switch. The connection joins the interface coming from the circuit's source to the interface leading to the circuit's destination. When you specify the interface names, include the logical portion of the name, which corresponds to the logical unit number. The cross-connect is bidirectional, so packets received on the first interface are transmitted from the second interface, and those received on the second interface are transmitted from the first.

To configure a connection for a local interface switch, include the following statements:

```
interface-switch connection-name {
 interface interface-name.unit-number;
}
lsp-switch connection-name {
 transmit-lsp lsp-number;
 receive-lsp lsp-number;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

To configure a connection for a remote interface switch, include the following statements:

```
remote-interface-switch connection-name {
 interface interface-name.unit-number;
 interface interface-name.unit-number;
 transmit-lsp lsp-number;
 receive-lsp lsp-number;
}
```

You can include these statements at the following hierarchy levels:

- [edit protocols connections]
- [edit logical-systems *logical-system-name* protocols connections]

## Configuring MPLS for Layer 2 Switching TCCs

For a Layer 2 switching TCC to work, you must enable MPLS on the router by including at least the following statements. This minimum configuration enables MPLS on a logical interface for the switching cross-connect.

Include the **family mpls** statement:

```
family mpls;
```

You can configure this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]

- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

You can then specify this logical interface in the MPLS protocol configuration:

```
mpls {
 interface interface-name; # Required to enable MPLS on the interface
}
```

You can configure these statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]



**NOTE:** MPLS LSP link protection does not support TCC.


## CCC and TCC Graceful Restart

CCC and TCC graceful restart allows Layer 2 connections between customer edge (CE) routers to restart gracefully. These Layer 2 connections are configured with the **remote-interface-switch** or **lsp-switch** statements. Because these CCC and TCC connections have an implicit dependency on RSVP LSPs, graceful restart for CCC and TCC uses the RSVP graceful restart capabilities.

RSVP graceful restart must be enabled on the PE routers and P routers to enable graceful restart for CCC and TCC. Also, because RSVP is used as the signaling protocol for signaling label information, the neighboring router must use helper mode to assist with the RSVP restart procedures.

Figure 71 on page 664 illustrates how graceful restart might work on a CCC connection between two CE routers.

**Figure 71: Remote Interface Switch Connecting Two CE Routers Using CCC**



PE Router A is the ingress for the transmit LSP from PE Router A to PE Router B and the egress for the receive LSP from PE Router B to PE Router A. With RSVP graceful restart enabled on all the PE and P routers, the following occurs when PE router A restarts:

- PE Router A preserves the forwarding state associated with the CCC routes (those from CCC to MPLS and from MPLS to CCC).
- Traffic flows without disruption from CE router to CE router.
- After the restart, PE Router A preserves the label for the LSP for which PE Router A is the egress (the receive LSP, for example). The transmit LSP from PE Router A to PE Router B can derive new label mappings, but should not cause any traffic disruption.

---

## Configuring CCC and TCC Graceful Restart

To enable CCC and TCC graceful restart, include the **graceful-restart** statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options]**
- **[edit logical-systems *logical-system-name* routing-options]**

The **graceful-restart** statement enables graceful restart for all protocols supporting this feature on the router. For more information about graceful restart, see the *Junos OS Routing Protocols Library for Routing Devices*.

CCC and TCC graceful restart depend on RSVP graceful restart. If you disable RSVP graceful restart, CCC and TCC graceful restart will not work. For more information about RSVP graceful restart, see [“RSVP Graceful Restart” on page 511](#) and [“Configuring RSVP Graceful Restart” on page 514](#).

---

## Configuring CCC Switching for Point-to-Multipoint LSPs

You can configure circuit cross-connect (CCC) between two circuits to switch traffic from interfaces to point-to-multipoint LSPs. This feature is useful for handling multicast or broadcast traffic (for example, a digital video stream).

To configure CCC switching for point-to-multipoint LSPs, you do the following:

- On the ingress provider edge (PE) router, you configure CCC to switch traffic from an incoming interface to a point-to-multipoint LSP.
- On the egress PE, you configure CCC to switch traffic from an incoming point-to-multipoint LSP to an outgoing interface.

The CCC connection for point-to-multipoint LSPs is unidirectional.

For more information about point-to-multipoint LSPs, see [“Point-to-Multipoint LSPs Overview” on page 281](#).

To configure a CCC connection for a point-to-multipoint LSP, complete the steps in the following sections:

- [Configuring the Point-to-Multipoint LSP Switch on Ingress PE Routers on page 666](#)
- [Configuring Local Receivers on a Point-to-Multipoint CCC LSP Switch on Ingress PE Routers on page 666](#)
- [Configuring the Point-to-Multipoint LSP Switch on Egress PE Routers on page 667](#)

## Configuring the Point-to-Multipoint LSP Switch on Ingress PE Routers

To configure the ingress PE router with a CCC switch for a point-to-multipoint LSP, include the **p2mp-transmit-switch** statement:

```
p2mp-transmit-switch switch-name {
 input-interface input-interface-name.unit-number;
 transmit-p2mp-lsp transmitting-lsp;
}
```

You can include the **p2mp-transmit-switch** statement at the following hierarchy levels:

- **[edit protocols connections]**
- **[edit logical-systems *logical-system-name* protocols connections]**

***switch-name*** specifies the name of the ingress CCC switch.

**input-interface *input-interface-name.unit-number*** specifies the name of the ingress interface.

**transmit-p2mp-lsp *transmitting-lsp*** specifies the name of the transmitting point-to-multipoint LSP.

## Configuring Local Receivers on a Point-to-Multipoint CCC LSP Switch on Ingress PE Routers

In addition to configuring an incoming CCC interface to a point-to-multipoint LSP on an ingress PE router, you can also configure CCC to switch traffic on an incoming CCC interface to one or more outgoing CCC interfaces by configuring output interfaces as local receivers.

To configure output interfaces, include the **output-interface** statement at the **[edit protocols connections p2mp-transmit-switch *p2mp-transmit-switch-name*]** hierarchy level.

```
[edit protocols connections]
p2mp-transmit-switch pc-ccc {
 input-interface fe-1/3/1.0;
 transmit-p2mp-lsp myp2mp;
 output-interface [fe-1/3/2.0 fe-1/3/3.0];
}
```

You can configure one or more output interfaces as local receivers on the ingress PE router using this statement.

Use the **show connections p2mp-transmit-switch (extensive | history | status)**, **show route ccc <interface-name> (detail | extensive)**, and **show route forwarding-table ccc**



`<interface-name> (detail | extensive)` commands to view details of the local receiving interfaces on the ingress PE router.

## Configuring the Point-to-Multipoint LSP Switch on Egress PE Routers

To configure the CCC switch for a point-to-multipoint LSP on the egress PE router, include the `p2mp-receive-switch` statement.

```
p2mp-receive-switch switch-name {
 output-interface [output-interface-name.unit-number];
 receive-p2mp-lsp receptive-lsp;
}
```

You can include this statement at the following hierarchy levels:

- `[edit protocols connections]`
- `[edit logical-systems logical-system-name protocols connections]`

*switch-name* specifies the name of the egress CCC switch.

`output-interface [ output-interface-name.unit-number ]` specifies the name of one or more egress interfaces.

`receive-p2mp-lsp receptive-lsp` specifies the name of the receptive point-to-multipoint LSP.



## PART 6

# Configuring GMPLS

- [GMPLS Overview on page 671](#)
- [Configuring GMPLS on page 677](#)
- [Using a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs over a Single RSVP LSP on page 723](#)



## CHAPTER 17

# GMPLS Overview

- [Introduction to GMPLS on page 671](#)
- [GMPLS Terms and Acronyms on page 672](#)
- [Supported GMPLS Standards on page 673](#)
- [GMPLS Operation on page 674](#)
- [GMPLS and OSPF on page 675](#)
- [GMPLS and CSPF on page 675](#)
- [GMPLS Features on page 676](#)

### Introduction to GMPLS

---

Traditional MPLS is designed to carry Layer 3 IP traffic using established IP-based paths and associating these paths with arbitrarily assigned labels. These labels can be configured explicitly by a network administrator, or can be dynamically assigned by means of a protocol such as LDP or RSVP.

GMPLS generalizes MPLS in that it defines labels for switching varying types of Layer 1, Layer 2, or Layer 3 traffic. GMPLS nodes can have links with one or more of the following switching capabilities:

- Fiber-switched capable (FSC)
- Lambda-switched capable (LSC)
- Time-division multiplexing (TDM) switched-capable (TSC)
- Packet-switched capable (PSC)

Label-switched paths (LSPs) must start and end on links with the same switching capability. For example, routers can establish packet-switched LSPs with other routers. The LSPs might be carried over a TDM-switched LSP between SONET add/drop multiplexers (ADMs), which in turn might be carried over a lambda-switched LSP.

The result of this extension of the MPLS protocol is an expansion in the number of devices that can participate in label switching. Lower-layer devices, such as OXCs and SONET ADMs, can now participate in GMPLS signaling and set up paths to transfer data. A router can participate in signaling optical paths across a transport network.

Two service models determine the visibility that a client node (a router, for example) has into the optical core or transport network. The first is through a user-to-network interface (UNI), which is often referred to as the overlay model. The second is known as the peer model. Juniper Networks supports both models.



**NOTE:** There is not necessarily a one-to-one correspondence between a physical interface and a GMPLS interface. If a GMPLS connection uses a nonchannelized physical connector, the GMPLS label can use the physical port ID. However, the label for channelized interfaces often is based on a channel or time slot. Consequently, it is best to refer to GMPLS labels as identifiers for a resource on a traffic engineering link.

To establish LSPs, GMPLS uses the following mechanisms:

- An out-of-band control channel and a data channel—RSVP messages for LSP setup are sent over an out-of-band control network. Once the LSP setup is complete and the path is provisioned, the data channel is up and can be used to carry traffic. The Link Management Protocol (LMP) is used to define and manage the data channels between a pair of nodes. You can optionally use LMP to establish and maintain LMP control channels between peers running the same Junos OS Release.
- RSVP-TE extensions for GMPLS—RSVP-TE is already designed to signal the setup of packet LSPs. This has been extended for GMPLS to be able to request path setup for various kinds of LSPs (nonpacket) and request labels like wavelengths, time slots, and fibers as label objects.
- Bidirectional LSPs—Data can travel both ways between GMPLS devices over a single path, so nonpacket LSPs are signaled to be bidirectional.

---

## GMPLS Terms and Acronyms

### F

**Forwarding adjacency**     A forwarding path for sending data between GMPLS-enabled devices.

### G

**Generalized MPLS (GMPLS)**     An extension to MPLS that allows data from multiple layers to be switched over label-switched paths (LSPs). GMPLS LSP connections are possible between similar Layer 1, Layer 2, and Layer 3 devices.

**GMPLS label**     Layer 3 identifiers, fiber port, time-division multiplexing (TDM) time slot, or dense wavelength-division multiplexing (DWDM) wavelength of a GMPLS-enabled device used as a next-hop identifier.

- GMPLS LSP types**      The four types of GMPLS LSPs are:
- Fiber-switched capable (FSC)—LSPs are switched between two fiber-based devices, such as optical cross-connects (OXCs) that operate at the level of individual fibers.
  - Lambda-switched capable (LSC)—LSPs are switched between two DWDM devices, such as OXCs that operate at the level of individual wavelengths.
  - TDM-switched capable (TDM)—LSPs are switched between two TDM devices, such as SONET ADMs.
  - Packet-switched capable (PSC)—LSPs are switched between two packet-based devices, such as routers or ATM switches.

## L

- Link Management Protocol**      A protocol used to define a forwarding adjacency between peers and to maintain and allocate resources on the traffic engineering links.

## T

- Traffic engineering link**      A logical connection between GMPLS-enabled devices. Traffic engineering links can have addresses or IDs and are associated with certain resources or interfaces. They also have certain attributes (encoding-type, switching capability, bandwidth, and so on). The logical addresses can be routable, although this is not required because they are acting as link identifiers. Each traffic engineering link represents a forwarding adjacency between a pair of devices.

## Supported GMPLS Standards

Junos OS substantially supports the following RFCs and Internet drafts, which define standards for Generalized MPLS (GMPLS).

- RFC 3471, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description*

Only the following features are supported:

- Bidirectional LSPs (upstream label only)
- Control channel separation
- Generalized label (suggested label only)
- Generalized label request (bandwidth encoding only)
- RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*  
Only Section 9, "Fault Handling," is supported.
- RFC 4202, *Routing Extensions in Support of Generalized Multi-Protocol Label Switching*  
Only interface switching is supported.
- RFC 4206, *Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)*

- Internet draft draft-ietf-ccamp-gmpls-rsvp-te-ason-02.txt, *Generalized MPLS (GMPLS) RSVP-TE Signalling in support of Automatically Switched Optical Network (ASON)* (expires January 2005)

- Internet draft draft-ietf-ccamp-gmpls-sonet-sdh-08.txt, *Generalized Multi-Protocol Label Switching Extensions for SONET and SDH Control*

Only S,U,K,L,M-format labels and SONET traffic parameters are supported.

- Internet draft draft-ietf-ccamp-lmp-10.txt, *Link Management Protocol (LMP)*

- Internet draft draft-ietf-ccamp-ospf-gmpls-extensions-12.txt, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching*

The following sub-TLV types for the Link type, link, value (TLV) are not supported:

- Link Local/Remote Identifiers (type 11)
- Link Protection Type (type 14)
- Shared Risk Link Group (SRLG) (type 16)

The features described in Section 2 of the draft, “Implications on Graceful Restart,” are also not supported.

The Interface Switching Capability Descriptor (type 15) sub-TLV type is implemented, but only for packet switching.

- Internet draft draft-ietf-mpls-bundle-04.txt, *Link Bundling in MPLS Traffic Engineering*

#### **Related Documentation**

- [Supported LDP Standards on page 520](#)
- [Supported MPLS Standards on page 20](#)
- [Supported RSVP Standards on page 460](#)
- [Accessing Standards Documents on the Internet](#)

---

## **GMPLS Operation**

The basic functionality of GMPLS requires close interaction between RSVP and LMP. It works in the following sequence:

1. LMP notifies RSVP of the new entities:
  - Traffic engineering link (forwarding adjacency)
  - Resources available for the traffic engineering link
  - Control peer
2. GMPLS extracts the LSP attributes from the configuration and requests RSVP to signal one or more specific paths, which are specified by the traffic engineering link addresses.
3. RSVP determines the local traffic engineering link, corresponding control adjacency and active control channel, and transmission parameters (such as IP destination). It



requests that LMP allocate a resource from the traffic engineering link with the specified attributes. If LMP finds a resource matching the attributes, label allocation succeeds. RSVP sends a PathMsg hop by hop until it reaches the target router.

4. When the target router receives the PathMsg, RSVP again requests that LMP allocate a resource based on the signaled parameters. If label allocation succeeds, the router sends back a ResvMsg.
5. If the signaling is successful, a bidirectional optical path is provisioned.

## GMPLS and OSPF

---

You can configure OSPF for GMPLS. OSPF is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). OSPF uses link-state information to make routing decisions.

## GMPLS and CSPF

---

GMPLS introduces extra constraints for computing paths for GMPLS LSPs that use CSPF. These additional constraints affect the following link attributes:

- Signal type (minimum LSP bandwidth)
- Encoding type
- Switching type

These new constraints are populated in the traffic engineering database with the exchange of an interface-switching capability descriptor type, length, value (TLV) through an IGP.

The ignored constraints that are exchanged through the interface switching capability descriptor include:

- Maximum LSP bandwidth
- Maximum transmission unit (MTU)

The CSPF path computation is the same as in non-GMPLS environments, except that the links are also limited by GMPLS constraints.

Each link can have multiple interface-switching capability descriptors. All the descriptors are checked before a link is rejected.

The constraints are checked in the following order:

1. The signal type configured for the GMPLS LSP signifies the amount of bandwidth requested. If the desired bandwidth is less than the minimum LSP bandwidth, the interface-switching descriptor is rejected.
2. The encoding type of the link for the ingress and the egress interfaces should match. The encoding type is selected and stored at the ingress node after all the constraints are satisfied by the link and is used to select the link on the egress node.

3. The switching type of the links of the intermediate switches should match that of the GMPLS LSP specified in the configuration.

## GMPLS Features

---

The Junos OS includes the following GMPLS functionality:

- An out-of-band control plane makes it possible to signal LSP path setup.
- RSVP-TE extensions support additional objects beyond Layer 3 packets, such as ports, time slots, and wavelengths.
- The LMP protocol creates and maintains a database of traffic engineering links and peer information. Only the static version of this protocol is supported in the Junos OS. You can optionally configure LMP to establish and maintain LMP control channels between peers running the same Junos OS Release.
- Bidirectional LSPs are required between devices.
- Several GMPLS label types that are defined in RFC 3471, *Generalized MPLS—Signaling Functional Description*, such as MPLS, Generalized, SONET/SDH, Suggested, and Upstream, are supported. Generalized labels do not contain a type field, because the nodes should know from the context of their connection what type of label to expect.
- Traffic parameters facilitate GMPLS bandwidth encoding and SONET/SDH formatting.
- Other supported attributes include interface identification and errored interface identification, user-to-network (UNI)-style signaling, and secondary LSP paths.

## CHAPTER 18

# Configuring GMPLS

- [LMP Configuration Overview on page 677](#)
- [Configuring LMP Traffic Engineering Links on page 678](#)
- [Configuring LMP Peers on page 680](#)
- [Configuring RSVP and OSPF for LMP Peer Interfaces on page 685](#)
- [Configuring MPLS Paths for GMPLS on page 686](#)
- [Tracing LMP Traffic on page 687](#)
- [Configuring MPLS LSPs for GMPLS on page 688](#)
- [Gracefully Tearing Down GMPLS LSPs on page 690](#)
- [GMPLS RSVP-TE VLAN LSP Signaling Overview on page 692](#)
- [Example: Configuring GMPLS RSVP-TE VLAN LSP Signaling on page 698](#)

## LMP Configuration Overview

---

You need to configure the Link Management Protocol (LMP) to define the data channel connection and the control channel connection between devices. Include the following statements at the **[edit protocols link-management]** hierarchy level:

```
[edit protocols link-management]
peer peer-name {
 address address;
 control-channel control-channel-name;
 lmp-control-channel control-channel-interface {
 remote-address ip-address;
 }
 lmp-protocol {
 hello-dead-interval milliseconds;
 hello-interval milliseconds;
 retransmission-interval milliseconds;
 retry-limit number;
 passive;
 }
 te-link te-link-name;
}
te-link te-link-name {
 disable;
 interface interface-name {
 disable;
```

```

 local-address ip-address;
 remote-address ip-address;
 remote-id id-number;
 }
 label-switched-path lsp-name;
 local-address ip-address;
 remote-address ip-address;
 remote-id id-number;
}
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}

```



**NOTE:** Although you can include GMPLS configuration statements at the [edit logical-systems *logical-system-name*] hierarchy level, GMPLS is not supported on logical systems.

For information about configuring LMP, see the following sections:

- [Configuring LMP Traffic Engineering Links on page 678](#)
- [Configuring LMP Peers on page 680](#)
- [Configuring RSVP and OSPF for LMP Peer Interfaces on page 685](#)
- [Configuring MPLS Paths for GMPLS on page 686](#)
- [Tracing LMP Traffic on page 687](#)

## Configuring LMP Traffic Engineering Links

An LMP traffic engineering link acts as a data channel connection between GMPLS devices.

To configure a traffic engineering link, include the **te-link** statement at the [edit protocols link-management] hierarchy level:

```

[edit protocols link-management]
te-link te-link-name {
 disable;
 interface interface-name {
 local-address ip-address;
 remote-address ip-address;
 remote-id id-number;
 }
 label-switched-path lsp-name;
 local-address ip-address;
 remote-address ip-address;
 remote-id id-number;
}

```

Complete the procedures in the following sections to configure an LMP traffic engineering link:

- [Configuring the Local IP Address for Traffic Engineering Links on page 679](#)
- [Configuring the Remote IP Address for Traffic Engineering Links on page 679](#)
- [Configuring the Remote ID for Traffic Engineering Links on page 680](#)

When you configure a traffic engineering link that contains interfaces for an LMP peer, you must also configure a control channel. However, no control channel is required for a traffic engineering link that contains an LSP. For information about configuring control channels, see [“Configuring LMP Peers” on page 680](#).

## Configuring the Local IP Address for Traffic Engineering Links

Use the **local-address** statement to configure the local IP address associated with the traffic engineering link.

We recommend that you configure an IP address subnet for your traffic engineering link addresses that is different from the subnet configured for your physical interfaces. This configuration enables you to identify which addresses are physical and which addresses belong to the traffic engineering link.

To configure the local IP address for the traffic engineering link, include the **local-address** statement:

```
te-link te-link-name {
 interface interface-name {
 local-address ip-address;
 }
 local-address ip-address;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Configuring the Remote IP Address for Traffic Engineering Links

You need to specify the address of the remote end of the data channel for each traffic engineering link. Use the **remote-address** statement to configure the remote IP address.

We recommend that you configure an IP address subnet for your traffic engineering link addresses that is different from the subnet configured for your physical interfaces. This enables you to identify which addresses are physical and which addresses belong to the traffic engineering link.

To configure the remote IP address for the traffic engineering link, include the **remote-address** statement:

```
te-link te-link-name {
 interface interface-name {
 remote-address ip-address;
 }
 remote-address ip-address;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

## Configuring the Remote ID for Traffic Engineering Links

The local ID for the traffic engineering link is automatically assigned by LMP. The port identifier and labels for the interfaces (resources) in the traffic engineering link are also assigned automatically. However, you need to explicitly configure the remote ID for the traffic engineering link and the remote ID traffic engineering link interface. The remote ID for the interface must be based on the post-ID assignment of the peer node. The remote IDs are needed for static mapping of remote labels to local labels.

Before you can obtain the remote IDs for the traffic engineering link and traffic engineering link interface on the peer node, you must first configure the LMP peer, as described in [“Configuring LMP Peers” on page 680](#). Once you have configured the LMP peer, you can obtain the traffic engineering link local ID and interface local ID by issuing the **show link-management te-link** command. Once you have these IDs, you can configure them as the remote IDs on the peer node.

To configure the remote ID for a traffic engineering link and for the traffic engineering link interface, include the **remote-id** statement:

```
te-link te-link-name {
 interface interface-name {
 remote-id id-number;
 }
 remote-id id-number;
}
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.

---

## Configuring LMP Peers

You need to configure network peers for GMPLS. A peer is a network device that your router communicates with when setting up the control and data channels. The peer is often an optical cross-connect (OXC).

To configure an LMP peer name, include the **peer** statement at the **[edit protocols link-management]** hierarchy level:

```
peer peer-name {
 address ip-address;
 control-channel control-channel-interface;
 lmp-control-channel control-channel-interface {
 remote-address ip-address;
 }
 lmp-protocol {
 hello-dead-interval milliseconds;
 hello-interval milliseconds;
 retransmission-interval milliseconds;
 retry-limit number;
 passive;
 }
}
```

```
te-link te-link-name;
}
```

The following sections describe how to configure an LMP peer:

- [Configuring the ID for LMP Peers on page 681](#)
- [Configuring the Interface for Control Channels Between LMP Peers on page 681](#)
- [Configuring the LMP Control Channel Interface for the Peer on page 681](#)
- [Configuring the Remote IP Address for LMP Control Channels on page 682](#)
- [Configuring Hello Message Intervals for LMP Control Channels on page 683](#)
- [Controlling Message Exchange for LMP Control Channels on page 684](#)
- [Preventing the Local Peer from Initiating LMP Negotiation on page 684](#)
- [Associating Traffic Engineering Links with LMP Peers on page 684](#)
- [Disabling the Traffic Engineering Link for LMP Peers on page 685](#)

## Configuring the ID for LMP Peers

To configure the LMP peer ID, include the **address** statement at the **[edit protocols link-management peer *peer-name*]** hierarchy level. The default value for the LMP peer ID is the loopback address.

```
[edit protocols link-management peer peer-name]
address ip-address;
```

## Configuring the Interface for Control Channels Between LMP Peers

You must configure one or more control channels between the LMP peers. The control channels must travel across either a point-to-point link or a tunnel.

To configure the interface for the control channel, include the **control-channel** statement at the **[edit protocols link-management peer *peer-name*]** hierarchy level:

```
[edit protocols link-management peer peer-name]
control-channel [interface-names];
```

You can configure a generic routing encapsulation (GRE) interface (*gre-x/y/z*) for the control channel. This type of interface does not require a Tunnel PIC.



**NOTE:** You can configure GRE interfaces (*gre-x/y/z*) only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information, see the *Junos OS Network Interfaces Library for Routing Devices*.

## Configuring the LMP Control Channel Interface for the Peer

In an environment that uses LMP to establish and maintain an LMP control channel between peers, you can configure a number of attributes associated with LMP. To configure the interface to be associated with the LMP control channel for the peer, include the **lmp-control-channel** statement:

**lmp-control-channel** *control-channel-interface*;

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management **peer** *peer-name*]
- [edit logical-systems *logical-system-name* protocols link-management **peer** *peer-name*]

You can configure a GRE interface for the LMP control channel. This type of interface does not require a Tunnel PIC.



**NOTE:** You can configure GRE interfaces only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information, see the *Junos OS Network Interfaces Library for Routing Devices*.

When this LMP control channel interface comes up, the peers use LMP to negotiate channel parameters and configure the control channel.

The local peer repeatedly sends a Config message to the remote peer. The Config message contains the local control channel ID, the local peer's node ID, a message ID, and a CONFIG object that includes hello message attributes (the hello interval and the hello dead interval).

The channel is activated when the remote peer responds with a ConfigAck message. The remote peer does so only when its own configured hello interval and hello dead interval match the values in the received Config message or the default values. If these values do not match, the remote peer responds with a ConfigNack message. The local peer logs this event and resends the Config message until the message retry limit is reached. When the message retry limit is reached, the local peer logs that event and restarts the configuration process.

## Configuring the Remote IP Address for LMP Control Channels

You need to specify the address of the remote end of the LMP control channel.

To configure the remote IP address for the LMP control channel, include the **remote-address** statement:

**remote-address** *address*;

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management **peer** *peer-name* **lmp-control-channel** *control-channel-interface*]
- [edit logical-systems *logical-system-name* protocols link-management **peer** *peer-name* **lmp-control-channel** *control-channel-interface*]



## Configuring Hello Message Intervals for LMP Control Channels

Hello messages are exchanged between LMP peers to maintain the control channel after LMP has activated the control channel. The LMP control channel is considered to be up only when the hello negotiation is successful. Successful negotiation consists of the local peer sending a hello message to the remote peer and receiving a hello message in response.

The LMP peers continue to exchange hello messages after the LMP control channel is up in order to maintain the channel.

The hello interval specifies the interval between periodic hello messages. The hello dead interval specifies how long the local peer waits for a hello response before it declares the LMP control channel to be down. When the channel goes down, the local peer restarts the LMP control channel negotiation and configuration process.

You can specify a hello interval from 150 through 300,000 milliseconds. The default hello interval is 150 milliseconds.

You can specify a hello dead interval from 500 through 300,000 milliseconds. The default hello dead interval is 500 milliseconds.

To configure the attributes for hello messages exchanged between LMP peers, include the **hello-interval** and **hello-dead-interval** statements:

```
hello-dead-interval milliseconds;
hello-interval milliseconds;
```

You can configure these statements at the following hierarchy levels:

- [edit protocols link-management **peer peer-name lmp-protocol**]
- [edit logical-systems *logical-system-name* protocols link-management **peer peer-name lmp-protocol**]

When an LMP control channel comes up after a successful exchange of hello messages between LMP peers, LMP uses link property correlation to verify the traffic engineering and data link information on both sides of a link. To do so, the local peer sends a LinkSummary message for each traffic engineering link governed by the LMP control channel. The LinkSummary message contains information that characterizes the traffic engineering link and each data link in the traffic engineering link.

The local peer continues sending a LinkSummary message for each link until the remote peer responds with a LinkSummaryAck message or until the message retry limit is reached. When the message retry limit is reached, the local peer logs that event and restarts the link property correlation process.

When the remote peer receives a LinkSummary message, it examines its own link information. If this information agrees with that in the LinkSummary message, the remote peer responds with a LinkSummaryAck message. If the information is different, the remote peer responds with a LinkSummaryNack message.

## Controlling Message Exchange for LMP Control Channels

You can configure message attributes that control the exchange of LMP Config and LinkSummary messages. The retransmission interval specifies the interval between resubmitted LMP messages. The retry limit specifies how many times LMP sends a message before restarting the process.

You can specify a retransmission interval from 500 through 300,000 milliseconds. The default retransmission interval is 500 milliseconds.

You can specify a retry limit from 3 through 1000 attempts. The default number of retry attempts is three.

To configure attributes governing the exchange of LMP messages between peers, include the **retransmission-interval** and **retry-limit** statements:

```
retransmission-interval milliseconds;
retry-limit number;
```

You can configure these statements at the following hierarchy levels:

- [edit protocols link-management **peer** *peer-name* **lmp-protocol**]
- [edit logical-systems *logical-system-name* protocols link-management **peer** *peer-name* **lmp-protocol**]

## Preventing the Local Peer from Initiating LMP Negotiation

You can specify that the local peer does not initiate LMP negotiation. Instead, the local peer waits for the remote peer to configure the LMP control channel.

To configure the local peer to wait for the remote peer to configure the LMP control channel, include the **passive** statement:

```
passive;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management **peer** *peer-name* **lmp-protocol**]
- [edit logical-systems *logical-system-name* protocols link-management **peer** *peer-name* **lmp-protocol**]

## Associating Traffic Engineering Links with LMP Peers

To specify the name of a traffic engineering link to be associated with this peer, include the **te-link** statement at the [edit protocols link-management **peer** *peer-name*] hierarchy level:

```
[edit protocols link-management peer peer-name]
te-link te-link-name;
```

For information about how to configure a traffic engineering link, see “Configuring LMP Traffic Engineering Links” on page 678.

## Disabling the Traffic Engineering Link for LMP Peers

To disable a specific traffic engineering link, include the **disable** statement:

```
disable;
```

You can configure this statement at the following hierarchy levels:

- [edit protocols link-management **te-link** *te-link-name*]
- [edit logical-systems *logical-system-name* protocols link-management **te-link** *te-link-name*]

## Configuring RSVP and OSPF for LMP Peer Interfaces

After you have configured the LMP peers as described in “Configuring LMP Peers” on page 680, add the peer interfaces to RSVP and OSPF. The peer interface name must match the peer name configured in LMP. Once the peer interfaces are added to the protocols, the traffic engineering link local and remote addresses can be signaled and advertised to peers like any other interface enabled for RSVP and OSPF. These addresses act as virtual interfaces for GMPLS.



**NOTE:** When adding the virtual peer interfaces to RSVP and OSPF, do not configure the corresponding physical control channel interface in either protocol. If you include the **interface all** statement, you must disable RSVP and OSPF protocols manually on the control channel interface.

To configure peer interfaces in RSVP and OSPF, complete the procedures in the following sections:

- [Configuring RSVP Signaling for LMP Peer Interfaces on page 685](#)
- [Configuring OSPF Routing for LMP Peer Interfaces on page 686](#)
- [Configuring the Hello Interval for LMP Peer Interfaces on page 686](#)

## Configuring RSVP Signaling for LMP Peer Interfaces

To configure RSVP signaling for LMP peers, configure the LMP peer interface by including the **peer-interface** statement at the [edit protocols **rsvp**] hierarchy level:

```
[edit protocols rsvp]
peer-interface peer-interface-name {
 (aggregate | no-aggregate);
 authentication-key key;
 disable;
 hello-interval seconds;
 (reliable | no-reliable);
}
```

The statements configured at the `[edit protocols rsvp peer-interface peer-interface-name]` hierarchy level have the same functionality as the statements configured at the `[edit protocols rsvp interface interface-name]` hierarchy level.

## Configuring OSPF Routing for LMP Peer Interfaces

To configure OSPF routing for LMP peers, configure the name of the LMP peer by including the `peer-interface` statement at the `[edit protocols ospf area area-number]` hierarchy level:

```
[edit protocols ospf area area-number]
peer-interface peer-interface-name {
 dead-interval seconds;
 disable;
 hello-interval seconds;
 retransmit-interval seconds;
 transit-delay seconds;
}
```

For information about how to configure OSPF statements, see the *Junos OS Routing Protocols Library for Routing Devices*.

## Configuring the Hello Interval for LMP Peer Interfaces

Hello packets are used to indicate to neighboring routers that the peer interface is still up and running. The hello interval must be the same for all routers on a shared logical IP network. You can specify a hello interval from 1 through 255 seconds. The default hello interval is normally 10 seconds. For nonbroadcast networks, the default hello interval is 120 seconds.

To specify how often the router sends hello packets out the peer interface, configure the `hello-interval` statement:

```
hello-interval seconds;
```

You can configure this statement at the following hierarchy levels:

- `[edit protocols ospf area area-number peer-interface peer-interface-name]`
- `[edit logical-systems logical-system-name protocols ospf area area-number peer-interface peer-interface-name]`

---

## Configuring MPLS Paths for GMPLS

As part of the configuration for GMPLS, you need to establish an MPLS path for each unique device connected through GMPLS. Configure the traffic engineering link remote address as the address at the `[edit protocols mpls path path-name]` hierarchy level. Constrained Shortest Path First (CSPF) is supported so you can choose either the `strict` or `loose` option with the address.

See “[LMP Configuration Overview](#)” on page 677 for information about how to obtain a traffic engineering link remote address.

To configure the MPLS path, include the `path` statement at the `[edit protocols mpls]` hierarchy level:

```
[edit protocols mpls]
path path-name {
 next-hop-address (strict | loose);
}
```

For information about how to configure MPLS paths, see [“Creating Named Paths” on page 60](#).

## Tracing LMP Traffic

To trace LMP protocol traffic, include the **traceoptions** statement at the **[edit protocols link-management]** hierarchy level:

```
[edit protocols link-management]
traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag flag <flag-modifier> <disable>;
}
```

Use the **file** statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory `/var/log`.

The following trace flags display the operations associated with the sending and receiving of various LMP messages:

- **all**—Trace all available operations
- **hello-packets**—Trace hello packets on any LMP control channel
- **init**—Output from the initialization messages
- **packets**—Trace all packets other than hello packets on any LMP control channel
- **parse**—Operation of the parser
- **process**—Operation of the general configuration
- **route-socket**—Operation of route socket events
- **routing**—Operation of the routing protocols
- **server**—Server processing operations
- **show**—Servicing operations for **show** commands
- **state**—Trace state transitions of the LMP control channels and traffic engineering links

Each flag can carry one or more of the following flag modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

## Configuring MPLS LSPs for GMPLS

---

To enable the proper GMPLS switching parameters, configure the label-switched path (LSP) attributes that are appropriate for your network connection. The default value for **switching-type** is **psc-1**, which is also appropriate for standard MPLS.

To configure the LSP attributes, include the **lsp-attributes** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name]
lsp-attributes {
 encoding-type type;
 gp-id gp-id;
 signal-bandwidth type;
 switching-type type;
}
```

If you include the **no-cspf** statement in the label-switched path configuration, you must also configure primary and secondary paths, or the configuration cannot be committed.

The following sections describe how to configure each of the LSP attributes for a GMPLS LSP:

- [Configuring the Encoding Type on page 688](#)
- [Configuring the GPID on page 689](#)
- [Configuring the Signal Bandwidth Type on page 689](#)
- [Configuring GMPLS Bidirectional LSPs on page 689](#)
- [Allowing Nonpacket GMPLS LSPs to Establish Paths Through Routers Running Junos OS on page 690](#)

### Configuring the Encoding Type

You need to specify the encoding type of the payload carried by the LSP. It can be any of the following:

- **ethernet**—Ethernet
- **packet**—Packet
- **pdh**—Plesiochronous digital hierarchy (PDH)
- **sonet-sdh**—SONET/SDH

The default value is **packet**.

To configure the encoding type, include the **encoding-type** statement at the **[edit protocols mpls label-switched-path *lsp-name* lsp-attributes]** hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
encoding-type type;
```

## Configuring the GPID

You need to specify the type of payload carried by the LSP. The payload is the type of packet underneath the MPLS label. The payload is specified by the generalized payload identifier (GPID).

You can specify the GPID with any of the following values:

- **hdlc**—High-Level Data Link Control (HDLC)
- **ethernet**—Ethernet
- **ipv4**—IP version 4 (default)
- **pos-scrambling-crc-16**—For interoperability with other vendors' equipment
- **pos-no-scrambling-crc-16**—For interoperability with other vendors' equipment
- **pos-scrambling-crc-32**—For interoperability with other vendors' equipment
- **pos-no-scrambling-crc-32**—For interoperability with other vendors' equipment
- **ppp**—Point-to-Point Protocol (PPP)

To configure the GPID, include the **gpid** statement at the **[edit protocols mpls label-switched-path *lsp-name* lsp-attributes]** hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
 gpid gpid;
```

## Configuring the Signal Bandwidth Type

The signal bandwidth type is the encoding used for path computation and admission control. To configure the signal bandwidth type, include the **signal-bandwidth** statement at the **[edit protocols mpls label-switched-path *lsp-name* lsp-attributes]** hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
 signal-bandwidth type;
```

## Configuring GMPLS Bidirectional LSPs

Because MPLS and GMPLS use the same configuration hierarchy for LSPs, it is helpful to know which LSP attributes control LSP functionality. Standard MPLS packet-switched LSPs are unidirectional, whereas GMPLS nonpacket LSPs are bidirectional.

If you use the default packet-switching type of **psc-1**, your LSP becomes unidirectional. To enable a GMPLS bidirectional LSP, you must select a non-packet-switching type option, such as **lambda**, **fiber**, or **ethernet**. Include the **switching-type** statement at the **[edit protocols mpls label-switched-path *lsp-name* lsp-attributes]** hierarchy level:

```
[edit protocols mpls label-switched-path lsp-name lsp-attributes]
 switching-type (lambda | fiber | ethernet);
```

## Allowing Nonpacket GMPLS LSPs to Establish Paths Through Routers Running Junos OS

By setting the A-bit in the Admin Status object, you can enable nonpacket GMPLS LSPs to establish paths through routers that run Junos. When an ingress router sends an RSVP PATH message with the Admin Status A-bit set, an external device (not a router running the Junos OS) can either perform a Layer 1 path setup test or help bring up an optical cross-connect.

When set, the A-bit in the Admin Status object indicates the administrative down status for a GMPLS LSP. This feature is used specifically by nonpacket GMPLS LSPs. It does not affect control path setup or data forwarding for packet LSPs.

Junos does not distinguish between the control path setup and data path setup. Other nodes along the network path use RSVP PATH signaling using the A-bit in a meaningful way.

To configure the Admin Status object for a GMPLS LSP, include the **admin-down** statement:

**admin-down;**

You can include this statement at the following hierarchy levels:

- **[edit protocols mpls label-switched-path *lsp-name*]**
- **[edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*]**

---

## Gracefully Tearing Down GMPLS LSPs

You can gracefully tear down nonpacket GMPLS LSPs. An LSP that is torn down abruptly, a common process in a packet-switched network, can cause stability problems in nonpacket-switched networks. To maintain the stability of nonpacket-switched networks, it might be necessary to tear down LSPs gracefully.

The following sections describe how to tear down GMPLS LSPs gracefully:

- [Temporarily Deleting GMPLS LSPs on page 690](#)
- [Permanently Deleting GMPLS LSPs on page 691](#)
- [Configuring the Graceful Deletion Timeout Interval on page 691](#)

### Temporarily Deleting GMPLS LSPs

You can gracefully tear down a GMPLS LSP using the **clear rsvp session gracefully** command.

This command gracefully tears down an RSVP session for a nonpacket LSP in two passes. In the first pass, the Admin Status object is signaled along the path to the endpoint of the LSP. During the second pass, the LSP is taken down. Using this command, the LSP is taken down temporarily. After the appropriate interval, the GMPLS LSP is resignaled and then reestablished.

The **clear rsvp session gracefully** command has the following properties:



- It only works on the ingress and egress routers of an RSVP session. If used on a transit router, it has the same behavior as the **clear rsvp session** command.
- It only works for nonpacket LSPs. If used with packet LSPs, it has the same behavior as the **clear rsvp session** command.

For more information, see the [CLI Explorer](#).

## Permanently Deleting GMPLS LSPs

When you disable an LSP in the configuration, the LSP is permanently deleted. By configuring the **disable** statement, you can disable a GMPLS LSP permanently. If the LSP being disabled is a nonpacket LSP, then the graceful LSP tear-down procedures that use the Admin Status object are used. If the LSP being disabled is a packet LSP, then the regular signaling procedures for LSP deletion are used.

To disable a GMPLS LSP, include the **disable** statement at any of the following hierarchy levels:

- **[edit protocols mpls label-switched-path *lsp-name*]**—Disable the LSP.
- **[edit protocols link-management *te-link te-link-name*]**—Disable a traffic engineering link.
- **[edit protocols link-management *te-link te-link-name interface interface-name*]**—Disable an interface used by a traffic engineering link.

## Configuring the Graceful Deletion Timeout Interval

The router that initiates the graceful deletion procedure for an RSVP session waits for the graceful deletion timeout interval to ensure that all routers along the path (especially the ingress and egress routers) have prepared for the LSP to be taken down.

The ingress router initiates the graceful deletion procedure by sending the Admin Status object in the path message with the **D** bit set. The ingress router expects to receive an Resv message with the **D** bit set from the egress router. If the ingress router does not receive this message within the time specified by the graceful deletion timeout interval, it initiates a forced tear-down of the LSP by sending a PathTear message.

To configure the graceful deletion timeout interval, include the **graceful-deletion-timeout** statement at the **[edit protocols rsvp]** hierarchy level. You can configure a time between 1 through 300 seconds. The default value is 30 seconds.

**graceful-deletion-timeout** *seconds*;

You can configure this statement at the following hierarchy levels:

- **[edit protocols rsvp]**
- **[edit logical-systems *logical-system-name* protocols rsvp]**

You can use the **show rsvp version** command to determine the current value configured for the graceful deletion timeout.

## GMPLS RSVP-TE VLAN LSP Signaling Overview

---

- [Understanding GMPLS RSVP-TE Signaling on page 692](#)
- [Need for GMPLS RSVP-TE VLAN LSP Signaling on page 692](#)
- [GMPLS RSVP-TE VLAN LSP Signaling Functionality on page 694](#)
- [LSP Hierarchy with GMPLS RSVP-TE VLAN LSP on page 695](#)
- [Path Specification for GMPLS RSVP-TE VLAN LSP on page 695](#)
- [GMPLS RSVP-TE VLAN LSP Configuration on page 695](#)
- [Associated Bidirectional Packet LSP on page 696](#)
- [Make-Before-Break for Associated Bidirectional Packet and GMPLS RSVP-TE VLAN LSP on page 697](#)
- [Supported and Unsupported Features on page 698](#)

### Understanding GMPLS RSVP-TE Signaling

Signaling is the process of exchanging messages within the control plane to set up, maintain, modify, and terminate data paths (label-switched paths (LSPs)) in the data plane. Generalized MPLS (GMPLS) is a protocol suite that extends the existing control plane of MPLS to manage further classes of interfaces and to support other forms of label switching, such as time-division multiplexing (TDM), fiber (port), Lambda, and so on.

GMPLS extends intelligent IP/MPLS connections from Layer 2 and Layer 3 all the way to Layer 1 optical devices. Unlike MPLS, which is supported mainly by routers and switches, GMPLS can also be supported by optical platforms, including SONET/SDH, optical cross-connects (OXCs), and dense wave division multiplexing (DWDM).

In addition to labels, which are primarily used to forward data in MPLS, other physical entries, such as wavelengths, time slots, and fibers can be used as label objects to forward data in GMPLS, thereby leveraging the existing control plane mechanisms to signal different kinds of LSPs. GMPLS uses RSVP-TE to be able to request the other label objects to signal the various kinds of LSPs (nonpacket). Bidirectional LSPs and an out-of-band control channel and a data channel using the Link Management Protocol (LMP) are the other mechanisms that are used by GMPLS to establish LSPs.


### Need for GMPLS RSVP-TE VLAN LSP Signaling

The traditional Layer 2 point-to-point services use Layer 2 circuits and Layer 2 VPN technologies that are based on LDP and BGP. In the traditional deployment, the customer edge (CE) devices do not participate in the signaling of the Layer 2 service. The provider edge (PE) devices manage and provision the Layer 2 service to provide end-to-end connectivity between the CE devices.

One of the biggest challenges of having the PE devices provision the Layer 2 services for each Layer 2 circuit between a pair of CE devices is the network management burden on the provider network.

Figure 72 on page 693 illustrates how the Layer 2 service is set up and used by the CE routers in a LDP/BGP-based Layer 2 VPN technology. Two CE routers CE1 and CE2 are connected to a provider MPLS network through the PE routers PE1 and PE2 respectively. The CE routers are connected to the PE routers by Ethernet links. Routers CE1 and CE2 are configured with VLAN1 and VLAN2 logical Layer 3 interfaces, so they appear to be directly connected. Routers PE1 and PE2 are configured with Layer 2 circuit (pseudowire) to carry the Layer 2 VLAN traffic between the CE routers. The PE routers use packet MPLS LSPs within the provider MPLS network to carry the Layer 2 VLAN traffic.

**Figure 72: Traditional Layer 2 Point-to-Point Services**




With the introduction of GMPLS-based VLAN LSP signaling, the need for the PE (also called server-layer) network to provision each individual Layer 2 connection between the CE (also called client) devices is minimized. The client router requests the server-layer router to which it is directly connected, for setting up the Layer 2 service to connect with a remote client router through GMPLS signaling.

The server-layer devices extend the signaling through the server-layer network to connect with the remote client routers. In the process, the server-layer device sets up the data plane for the Layer 2 service at the server-client border, and sets up the data plane for carrying the Layer 2 traffic within the server-layer network. With the Layer 2 service setup, the client routers can run IP/MPLS directly on top of the Layer 2 service and have IP/MPLS adjacency with each other.

In addition to reducing the provisioning activity needed on the server-layer devices, GMPLS signaling also provides the client routers with the flexibility of bringing up the Layer 2 circuits on an on-demand basis without depending on the server-layer administration for the provisioning of the Layer 2 service.

Using the same topology as in Figure 1, Figure 73 on page 694 illustrates how the Layer 2 service is set up and used by the client routers in GMPLS RSVP-TE-based Layer 2 VPN technology.

Figure 73: GMPLS RSVP-TE VLAN LSP



In [Figure 73 on page 694](#), instead of configuring a pseudowire to carry the Layer 2 VLAN traffic between the client routers, Routers PE1 and PE2 are configured with an IP-based communication channel and other GMPLS-specific configurations (identification of Ethernet links as TE-links) for allowing the exchange of GMPLS RSVP-TE signaling messages with the client routers. Routers CE1 and CE2 are also configured with an IP-based communication channel and relevant GMPLS configuration for exchanging the GMPLS RSVP-TE signaling messages with the server-layer routers. Routers CE1 and CE2 establish an IP/MPLS adjacency on top of this Layer 2 service.

### GMPLS RSVP-TE VLAN LSP Signaling Functionality

Based on [Figure 73 on page 694](#), the client router establishes the Layer 2 service in the server-layer network as follows:

1. Router CE1 initiates GMPLS RSVP-TE signaling with Router PE1. In this signaling message, Router CE1 indicates the VLAN on the Ethernet link for which it needs the Layer 2 service and the remote CE router, Router CE2, with which the VLAN should be connected.  
  
Router CE1 also indicates the remote PE router, Router PE2, to which Router CE2 is connected, and the exact Ethernet link connecting Router CE2 to Router PE2 on which the Layer 2 service is required in the signaling message.
2. Router PE1 uses the information from Router CE1 in the signaling message and determines the remote PE router, Router PE2, with which Router CE2 is attached. Router PE1 then establishes a packet MPLS LSP (associated bidirectional) through the server-layer MPLS network for carrying the VLAN traffic and then passes the GMPLS RSVP-TE signaling message to Router PE2 using the LSP hierarchy mechanism.
3. Router PE2 propagates the GMPLS RSVP-TE signaling message to Router CE2 with the VLAN to be used on the PE2-CE2 Ethernet link.
4. Router CE2 responds with an acknowledgment to the GMPLS RSVP-TE signaling message to Router PE2. Router PE2 then propagates it to Router PE1, which in turn propagates it to Router CE1.
5. As part of this message propagation, Routers PE1 and PE2 set up the forwarding plane to enable bidirectional flow of VLAN Layer 2 traffic between Routers CE1 and CE2.

## LSP Hierarchy with GMPLS RSVP-TE VLAN LSP

The Layer 2 service in GMPLS RSVP-TE VLAN LSP signaling is brought up using a hierarchy mechanism in which two different RSVP LSPs are created for the Layer 2 service:

- An end-to-end VLAN LSP that has state information at the client and server-layer routers.
- An associated bidirectional packet transport LSP that is present in the server-layer routers (PE and P) of the server-layer network.

The LSP hierarchy avoids sharing information about technology-specific LSP characteristics with the core nodes of the server-layer network. This solution cleanly separates the VLAN LSP state and the transport LSP state, and ensures that the VLAN LSP state is only present on the nodes (PE, CE) where it is needed.

## Path Specification for GMPLS RSVP-TE VLAN LSP

The path for the GMPLS RSVP-TE LSP is configured as an Explicit Route Object (ERO) at the initiating client router. As this LSP traverses different network domains (initiating, terminating at client network, and traversing the server-layer network), the LSP setup falls under the category of an interdomain LSP setup. In an interdomain scenario, one network domain generally does not have full visibility into the topology of the other network domain. Hence, the ERO that gets configured at the initiating client router does not have full hop information for the server-layer portion. This feature requires that the ERO configured at the CE router has three hops, with the first hop being a strict hop identifying the CE1-PE1 Ethernet link, the second hop being a loose hop identifying the egress PE router (PE2), and the third hop being a strict hop identifying the CE2-PE2 Ethernet link.

## GMPLS RSVP-TE VLAN LSP Configuration

The configuration required to set up a GMPLS VLAN LSP at the client and server routers uses the existing GMPLS configuration model with some extensions. The Junos OS GMPLS configuration model for nonpacket LSPs is targeted toward bringing the physical interfaces up and running through GMPLS RSVP-TE signaling, whereas signaling a GMPLS RSVP-TE VLAN LSP aims at bringing up individual VLANs on top of a physical interface. The **ethernet-vlan** configuration statement under the **[edit protocols link-management te-link]** hierarchy enables this.

The client router has physical interfaces connected to a server network, and the server network provides a point-to-point connection between two client routers over the attached physical interfaces. The physical interface is brought into an operational state by GMPLS RSVP-TE as follows:

1. The client router maintains a routing or signaling adjacency with the server network node to which the physical interface is connected, typically through a control channel different from the physical interface, because the physical interface itself is brought up and running only after the signaling.
2. The client router and the server network node identify the physical interfaces connecting them using the TE-link mechanism.

3. The client router and the server network node use the TE-link identifier (IP address) as the GMPLS RSVP hop and the physical interface identifier as the GMPLS label values in the GMPLS RSVP-TE signaling messages to bring the physical interface into an operational state.

In the existing GMPLS configuration, the server and client network nodes use the **protocols link-management peer *peer-name*** configuration statement to specify the adjacent peer node. Because a client router can have one or more physical interfaces connected to the server network node, these physical interfaces are grouped and identified by an IP address through the **protocols link-management te-link *link-name*** configuration statement. The TE-link is assigned a local IP address, a remote IP address, and a list of physical interfaces. The TE-link is then associated with the **protocols link-management peer *peer-name* te-link *te-link-list*** configuration statement.

The out-of-band control channel that is required for exchanging signaling messages is specified using the **protocols link-management peer *peer-name* control-channel *interface-name*** configuration statement. The existence of the server or client network node is made visible to the RSVP and IGP (OSPF) protocols through the **peer-interface *interface-name*** configuration statement under the **[edit protocols *rsvp*]** and **[edit protocols *ospf*]** hierarchy levels.

In the existing GMPLS configuration, the label (upstream label and resv label) that is carried in the signaling message is an integer identifier that identifies the physical interface that is required to be brought up. As the label is used to identify the physical interface, the existing GMPLS configuration allows multiple interfaces to be grouped under a single TE-link. In the existing GMPLS configuration, there is sufficient information in the GMPLS RSVP-TE signaling message, such as TE-link address and label value, to identify the physical interface that is required to be brought up. In contrast, for GMPLS RSVP-TE VLAN LSP configuration, the VLAN ID value is used as the label in the signaling message.

In the GMPLS RSVP-TE VLAN LSP configuration, if multiple interfaces are allowed to be configured under a single TE-link, using VLAN ID as the label value in the signaling message can cause ambiguity as to which physical interface on which the VLAN has to be provisioned. Therefore, the TE-link is configured with the **ethernet-vlan** configuration statement, if the number of physical interfaces that can be configured under the TE-link is restricted to only one.

In the existing GMPLS configuration, the bandwidth for a nonpacket LSP is a discrete quantity that corresponds to the bandwidth of the physical interface that needs to be brought up. So, the GMPLS LSP configuration does not allow any bandwidth to be specified, but allows the bandwidth to be specified only through the **signal-bandwidth** configuration statement under the **[protocols mpls label-switched-path *lsp-name* lsp-attributes]** hierarchy level. In the GMPLS VLAN LSP configuration, bandwidth is specified similar to that of a packet LSP. In the GMPLS VLAN LSP configuration, the **bandwidth** option is supported and **signal-bandwidth** is not supported.

## Associated Bidirectional Packet LSP

The GMPLS RSVP-TE VLAN LSP is carried on an associated bidirectional transport LSP within the server-layer network, which is a single-sided provisioned LSP. The transport LSP signaling is initiated as a unidirectional LSP from the source router to the destination

router in the forward direction, and the destination router in turn initiates the signaling of the unidirectional LSP in the reverse direction back to the source router.

### Make-Before-Break for Associated Bidirectional Packet and GMPLS RSVP-TE VLAN LSP

The make-before-break support for an associated bidirectional transport LSP follows a similar model, where the destination router for the forward direction of the bidirectional LSP does not perform any make-before-break operations on the reverse direction of the bidirectional LSP. It is the source router (initiator of the associated bidirectional LSP) that initiates the make-before-break newer instance of the associated bidirectional LSP, and the destination router in turn initiates the make-before-break newer instance in the other direction.

For instance, in [Figure 73 on page 694](#), the unidirectional transport LSP is initiated from Router PE1 to Router PE2 in the forwarding direction, and in turn Router PE2 initiates the transport LSP to Router PE1 in the reverse direction. When a make-before-break instance occurs, only Router PE1 as the initiating client router can establish a new instance of the associated bidirectional LSP. Router PE2 in turn initiates the make-before-break newer instance in the reverse direction.

The make-before-break support for the associated bidirectional transport LSP is used only in scenarios where the transport LSP gets into a state of being locally protected due to link or node failure on the path of the LSP. The GMPLS RSVP-TE VLAN LSP uses the make-before-break mechanism for adjusting seamless bandwidth changes.



**NOTE:** Periodic re-optimization is not enabled for the associated bidirectional transport LSPs.

The newer make-before-break instance of the GMPLS VLAN LSP is supported under the following constraints:

- It should originate from the same client router as the older instance and be destined to the same client router as the older instance.
- It should use the same server-client links at both the server-client ends as the older instance.
- It should use the same VLAN label at the server-client links as the older instance.
- The GMPLS VLAN LSP should be configured as **adaptive** when the bandwidth change is initiated from the CLI, or else the current instance of the VLAN LSP is torn down and a new VLAN LSP instance is established.

The make-before-break operation for the GMPLS VLAN LSP on the server-layer edge router is rejected if these constraints are not met.

On the server-layer edge routers, when a make-before-break instance of the GMPLS VLAN LSP is seen, a completely new, separate associated bidirectional transport LSP is created to support this make-before-break instance. The existing associated bidirectional LSP (supporting the older instance) is not triggered to initiate a make-before-break instance at the transport LSP level. An implication of this choice (of initiating a new

transport LSP) is that at the server-layer resource/bandwidth sharing does not happen when a make-before-break operation is performed for the GMPLS VLAN LSP.

## Supported and Unsupported Features

Junos OS supports the following features with the GMPLS RSVP-TE VLAN LSP:

- Request for specific bandwidth and local protection for the VLAN LSP on the client router to the server-layer router.
- Nonstop active routing (NSR) support for the GMPLS VLAN LSP at the client routers, server-layer edge routers, and associated bidirectional transport LSP at the server-layer edge routers.
- Multichassis support.

Junos OS does **not** support the following GMPLS RSVP-TE VLAN LSP functionality:

- Graceful restart support for associated bidirectional packet LSP and GMPLS VLAN LSP.
- End-to-end path computation for GMPLS VLAN LSP using CSPF algorithm at the client router.
- Non-CSPF routing-based discovery of next-hop routers by the different client, server-layer edge routers.
- Automatic provisioning of the client Layer 3 VLAN interfaces upon the successful setup of the VLAN LSP at the client routers.
- MPLS OAM (LSP-ping, BFD).
- Packet MPLS applications, such as next-hop in static route and in IGP shortcuts.
- Local cross connect mechanism, where a client router connects to a remote client router which is connected to the same server router.
- Junos OS Services Framework.
- IPv6 support.
- Logical systems.
- Aggregated Ethernet/SONET/IRB interfaces at the server-client link.

### Related Documentation

- [Example: Configuring GMPLS RSVP-TE VLAN LSP Signaling on page 698](#)

---

## Example: Configuring GMPLS RSVP-TE VLAN LSP Signaling

This example shows how to configure GMPLS RSVP-TE VLAN LSP signaling on the client routers to enable one client router to connect with a remote client router through a server-layer network using the LSP hierarchy. This enables the client routers to establish, maintain, and provision the Layer 2 services, without depending on the server-layer



administration, thereby reducing the burden on the operational expenses of the provider network.

- [Requirements on page 699](#)
- [Overview on page 699](#)
- [Configuration on page 705](#)
- [Verification on page 715](#)

## Requirements

This example uses the following hardware and software components:

- Six routers that can be a combination of M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, T Series Core Routers, and PTX Series Packet Transport Routers
- Junos OS Release 14.2 or later running on the client routers and server-layer edge routers

Before you begin:

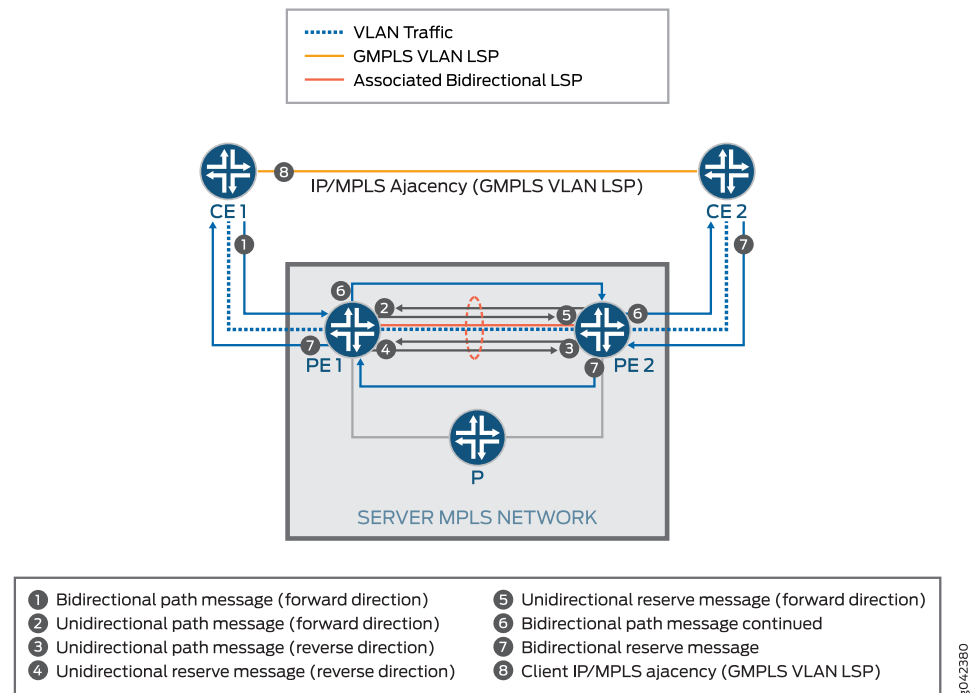
1. Configure the device interfaces.
2. Configure the interface-associated VLANs.
3. Configure the following routing protocols:
  - RSVP
  - MPLS
  - LMP

## Overview

Starting with Junos OS Release 14.2, the Layer 2 services between two client routers across an external/third-party server-layer network are set up by the client routers on an on-demand basis through GMPLS RSVP-TE signaling. This feature provides the client routers the flexibility to establish, maintain, and provision the Layer 2 services, without depending on the server-layer administration, thereby reducing the burden on the operational expenses of the provider network. In traditional Layer 2 VPN technology based on LDP and BGP, the provider network handled the provisioning activity for each Layer 2 circuit established between two client routers.

[Figure 74 on page 700](#) illustrates the setting up and signaling of the GMPLS VLAN LSP between two client routers, CE1 and CE2, across a server-layer network with two server-layer edge routers, PE1 and PE2, and one server-layer core router, P.

Figure 74: Setting Up a GMPLS VLAN LSP



The signaling of GMPLS VLAN LSP is executed as follows:

#### 1. Initiating GMPLS VLAN LSP at CE1

Router CE1 initiates the GMPLS VLAN LSP setup by sending the GMPLS RSVP-TE path message to Router PE1. The signaling between CE1 and PE1 is over an out-of-band control channel, which is a separate control VLAN configured on the Ethernet link connecting the two routers.

The GMPLS RSVP-TE path message initiated by Router CE1 is used to perform the following:

- Identify the Ethernet link on which the VLAN is active.
- Abstract the Ethernet link as a TE-link and assign an IP address to identify the Ethernet link.
- Allocate a VLAN ID from the pool of free VLANs managed by Router CE1 for every Ethernet link connecting Router PE1 to the identified Ethernet link.

This VLAN ID can also be used for the GMPLS VLAN LSP at the CE2-PE2 Ethernet link.

- Identify the VLAN for which the Layer 2 service is required to be set up using the allocated VLAN ID as the upstream label object and the upstream direction label value.

- e. Include an ERO object that helps Router PE1 in establishing the VLAN LSP through the server-layer network to the remote client router, CE2. The ERO object in the path message includes three hops:
  - First hop—Strict hop identifying the initiating client-server Ethernet link, PE1-CE1.
  - Second hop—Loose hop identifying the remote server-layer router, PE2.
  - Third hop—Strict hop identifying the remote client-server Ethernet link, PE2-CE2.
- f. Include the bandwidth required for the GMPLS VLAN LSP.
- g. Include any local-protection required within the server-layer network for the VLAN LSP.

## 2. Initiating Associated Bidirectional Transport LSP at PE1

After Router PE1 receives the path message from Router CE1, the message is validated to check the availability of the Ethernet link and VLAN ID. In the server-layer network, the Layer 2 services between the server-layer routers, PE1 and PE2, are provided at the data plane in a manner similar to Layer 2 circuits. Router PE1 brings up a transport LSP to Router PE2 and then extends the GMPLS VLAN LSP as a hierarchical LSP running on top of the PE1-PE2 transport LSP. The PE1-PE2 transport LSP is a packet LSP and is bidirectional in nature. This is because the GMPLS VLAN LSP is bidirectional and each server-layer router needs to be able to do the following:

- Receive traffic from the server-client Ethernet link (for example, the PE1-CE1 link) and send it to the remote server-layer router, PE2.
- Receive traffic from remote Router PE2 and send it on the PE1-CE1 Ethernet link.

For each GMPLS VLAN LSP, a packet transport LSP is set up within the server-layer network. The transport LSP is exclusively used to carry traffic of the GMPLS VLAN LSP for which it was created. The transport LSP is dynamically created at the time of receiving the GMPLS VLAN LSP; thus, no configuration is required to trigger its creation. The transport LSP established for the VLAN LSP inherits the bandwidth and the local-protection attributes from the VLAN LSP.

Router PE1 signals the PE1-PE2 transport LSP to Router PE2. Router PE1 determines the destination for the transport LSP from the loose hop specified in the ERO object of the GMPLS RSVP-TE path message from Router CE1 and then signals the VLAN LSP. However, if the PE1-PE2 transport LSP fails to establish, Router PE1 sends back a path error message to Router CE1, and the GMPLS VLAN LSP is not established as well.

## 3. Setting Up the Associated Bidirectional Transport LSP Between the Server-Layer Routers

The associated bidirectional LSP between routers PE1 and PE2 consists of two unidirectional packet LSPs:

- PE1-to-PE2
- PE2-to-PE1

Router PE1 initiates signaling of a unidirectional packet LSP to Router PE2. This unidirectional packet LSP constitutes the forward direction (PE1-to-PE2) of the

associated bidirectional LSP, and the path message carries the Extended Association Object indicating this is a single-sided provisioning model. On receiving the path message for the LSP, Router PE2 responds with a Resv message and triggers the signaling of a unidirectional packet LSP to Router PE1 with the same path as (PE1-to-PE2) in the reverse direction. This unidirectional packet LSP uses the PE2-to-PE1 direction of the associated bidirectional LSP, and this path message carries the same Extended Association Object seen in the PE1-to-PE2 path message.

When Router PE1 receives the Resv message for the PE1-to-PE2 unidirectional LSP and the path message for the PE2-to-PE1 unidirectional LSP, PE1 binds the PE1-to-PE2 and PE2-to-PE1 unidirectional LSPs by matching the Extended Association Objects carried in the respective path messages. For the path message for the PE2-to-PE1 unidirectional LSP, Router PE1 responds with the Resv Message. On receiving the Resv message for the PE1-to-PE2 LSP and the path message for the PE2-to-PE1 LSP, Router PE1 has established the associated bidirectional packet transport LSP.

#### **4. Setting Up the GMPLS VLAN LSP at Router PE1**

After successfully establishing the transport LSP, Router PE1 triggers the signaling of the GMPLS VLAN LSP. Router PE1 sends the GMPLS RSVP-TE path message corresponding to the VLAN LSP directly to Router PE2, which is bidirectional in nature and includes the upstream label object.

Router PE2 is not aware of the association between the transport LSP and the VLAN LSP. This association is indicated to Router PE2 by Router PE1.

#### **5. Setting Up the GMPLS VLAN LSP at Router PE2**

On receiving the VLAN LSP path message from Router PE1, Router PE2 verifies the availability of the transport LSP. If the transport LSP is not available or the LSP setup is in progress, the VLAN LSP processing is put on hold. When the transport LSP is available, Router PE2 processes the VLAN LSP path message. The ERO object in this path message indicates that the next hop is a strict hop identifying the PE2-to-CE2 Ethernet link. The ERO object can indicate the VLAN ID to be used on the PE2-to-CE2 Ethernet link by Router PE2.

Router PE2 appropriately allocates the VLAN ID to be sent as the upstream label in the VLAN LSP path message to Router CE2, and sends it through an out-of-band control channel.

#### **6. Processing the GMPLS VLAN LSP at Router CE2**

On receiving the GMPLS RSVP-TE LSP from Router PE2, Router CE2 validates the availability of VLAN ID for allocation on the PE2-to-CE2 link. Router CE2 then allocates the VLAN ID for this VLAN LSP and sends back a Resv message to Router PE2 with the VLAN ID as the label object in the Resv message.

#### **7. Processing the GMPLS VLAN LSP at Router PE2**

On receiving the Resv message from Router CE2, Router PE2 validates that the label object in the Resv message has the same VLAN ID as in the path message. Router PE2 then allocates a 20-bit MPLS label, which is included in the Resv message sent to Router PE1.

Router PE2 then programs the forwarding plane with the entries to provide the Layer 2 service functionality.



**NOTE:** For all the VLAN IDs that can be allocated as labels on the PE1-to-CE1 and PE2-CE2 Ethernet links, you must manually configure logical interfaces for circuit cross-connect (CCC) purposes on the server-layer edge routers and not for other families, such as IPv4, IPv6, or MPLS.

#### 8. Processing the GMPLS VLAN LSP at Router PE1

On receiving the Resv message for the VLAN LSP from Router PE2, Router PE1 sends a Resv message to Router CE1 with the same VLAN ID it received as the upstream label from Router CE1. Router PE1 programs the forwarding plane with the entries to provide the Layer 2 service functionality as Router PE2.

#### 9. Processing the GMPLS VLAN LSP at Router CE1

On receiving the Resv message from Router PE1, Router CE1 validates that the VLAN ID received in the Resv message matches the VLAN ID in the upstream label in the path message it sent. This completes the setup of the GMPLS VLAN LSP from Router CE1 to Router CE2.



**NOTE:**

- The GMPLS VLAN LSP setup does not result in the addition of any forwarding plane entries at the client routers, CE1 and CE2. Only the server-layer routers, PE1 and PE2, add the forwarding plane entries for the GMPLS VLAN LSP.
- There is no routing information exchange between the client and the server-layer routers. The client and server-layer routers do not exchange their network topology information with each other.

#### 10. Accounting for Bandwidth of the GMPLS VLAN LSP


On successfully setting up the GMPLS VLAN LSP, both the client and server-layer routers reduce the amount of available bandwidth on the server-client Ethernet links by the bandwidth amount allocated for the GMPLS VLAN LSP. This bandwidth accounting information is used for admission control purposes when additional GMPLS VLAN LSPs are brought up on the server-client Ethernet links.

#### 11. Using GMPLS VLAN LSP by the Client Routers

After successfully setting up the GMPLS VLAN LSP, the client routers – CE1 and CE2 – need to be manually configured with the VLAN logical interface on top of the server-client Ethernet links with the signaled VLAN ID. This logical interface needs to be configured with the IP address and needs to be included in the IGP protocol. As a result of this configuration, Routers CE1 and CE2 establish IGP adjacency and exchange data traffic over the Layer 2 service established through the GMPLS signaling.

Figure 75 on page 704 illustrates the data traffic flow of the GMPLS VLAN LSP from Router CE1 to Router CE2 after the LSP setup is complete and the necessary CE1-to-CE2 IGP/MPLS adjacency has been established. The server-layer transport LSP originates from Router PE1, traverses a single server-layer core router, Router P, and reaches Router PE2. The server-layer transport LSP is shown as a penultimate-hop pop LSP, where Router P pops off the transport LSP label, and only the service label is present on the P-to-PE2 link.

Figure 75: Data Traffic Flow of GMPLS VLAN LSP

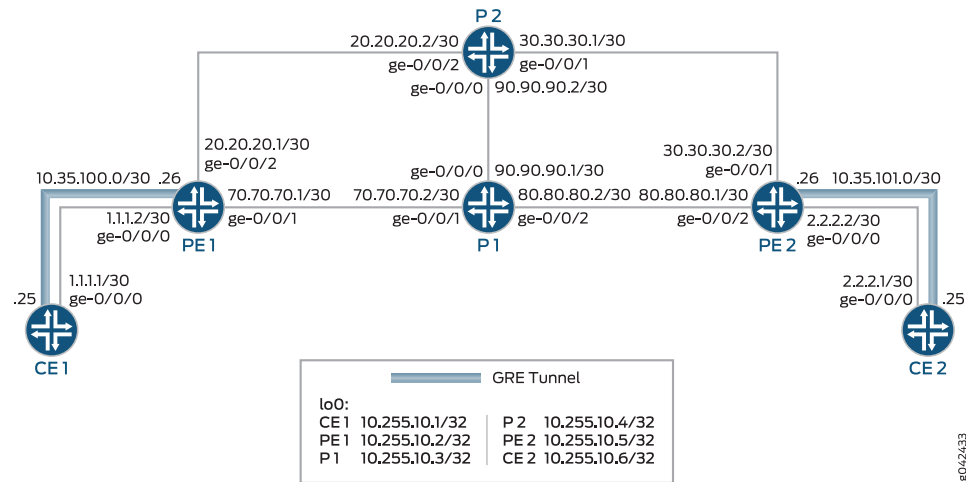


8042381

## Topology

In Figure 76 on page 705, GMPLS RSVP-TE VLAN LSP signaling is used to establish the Layer 2 services between the client routers, Router CE1 and Router CE2. The server routers, Router PE1 and Router PE2, have a GRE tunnel established with each of the directly connected client routers. Routers P1 and P2 are also server routers in the server-layer network.

Figure 76: Configuring GMPLS RSVP-TE VLAN LSP Signaling



## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
CE1 set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 1 vlan-id 1
set interfaces ge-0/0/0 unit 1 family inet address 1.1.1.1/30
set interfaces ge-0/0/0 unit 1 family mpls
set interfaces ge-0/0/0 unit 10 vlan-id 10
set interfaces ge-0/0/0 unit 10 family inet address 10.10.10.1/24
set interfaces ge-0/0/0 unit 10 family mpls
set interfaces gre unit 0 tunnel source 1.1.1.1
set interfaces gre unit 0 tunnel destination 1.1.1.2
set interfaces gre unit 0 family inet address 10.35.100.25/30
set interfaces lo0 unit 0 family inet address 10.255.10.1/32
set routing-options router-id 10.255.10.1
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols rsvp peer-interface PE1
set protocols mpls no-cspf
set protocols mpls label-switched-path CE1-to-CE2 from 10.255.10.1
set protocols mpls label-switched-path CE1-to-CE2 to 10.255.10.6
set protocols mpls label-switched-path CE1-to-CE2 lsp-attributes switching-type
 ethernet-vlan
set protocols mpls label-switched-path CE1-to-CE2 lsp-attributes upstream-label vlan-id
 10
set protocols mpls label-switched-path CE1-to-CE2 bandwidth 100m
set protocols mpls label-switched-path CE1-to-CE2 primary path1
set protocols mpls path path1 10.35.1.2 strict
set protocols mpls path path1 10.255.10.5 loose
set protocols mpls path path1 10.36.1.1 strict
set protocols mpls interface all
```

```

set protocols mpls interface fxp0.0 disable
set protocols link-management te-link link10 local-address 10.35.1.1
set protocols link-management te-link link10 remote-address 10.35.1.2
set protocols link-management te-link link10 ethernet-vlan
set protocols link-management te-link link10 interface ge-0/0/0
set protocols link-management peer PE1 address 10.255.10.2
set protocols link-management peer PE1 control-channel gre.0
set protocols link-management peer PE1 te-link link10

PE1 set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 encapsulation flexible-ethernet-services
set interfaces ge-0/0/0 unit 1 vlan-id 1
set interfaces ge-0/0/0 unit 1 family inet address 1.1.1.2/30
set interfaces ge-0/0/0 unit 1 family mpls
set interfaces ge-0/0/0 unit 10 encapsulation vlan-ccc
set interfaces ge-0/0/0 unit 10 vlan-id 10
set interfaces ge-0/0/1 unit 0 family inet address 70.70.70.1/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 20.20.20.1/30
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces gre unit 0 tunnel source 1.1.1.2
set interfaces gre unit 0 tunnel destination 1.1.1.1
set interfaces gre unit 0 family inet address 10.35.100.26/30
set interfaces lo0 unit 0 family inet address 10.255.10.2/32
set routing-options router-id 10.255.10.2
set protocols rsvp associated-bidirectional-lsp single-sided-provisioning
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols rsvp peer-interface CE1 dynamic-bidirectional-transport
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols link-management te-link link1 local-address 10.35.1.2
set protocols link-management te-link link1 remote-address 10.35.1.1
set protocols link-management te-link link1 ethernet-vlan vlan-id-range 1-1000
set protocols link-management te-link link1 interface ge-0/0/0
set protocols link-management peer CE1 address 10.255.10.1
set protocols link-management peer CE1 control-channel gre.0
set protocols link-management peer CE1 te-link link1

P1 set interfaces ge-0/0/0 unit 0 family inet address 90.90.90.1/24
set interfaces ge-0/0/0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family inet address 70.70.70.2/24
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces ge-0/0/2 unit 0 family inet address 80.80.80.2/24
set interfaces ge-0/0/2 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.10.3/32
set routing-options router-id 10.255.10.3
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols ospf traffic-engineering

```



```

set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable

P2 set interfaces ge-0/0/0 unit 0 family inet address 90.90.90.2/30
 set interfaces ge-0/0/0 unit 0 family mpls
 set interfaces ge-0/0/1 unit 0 family inet address 30.30.30.1/30
 set interfaces ge-0/0/1 unit 0 family mpls
 set interfaces ge-0/0/2 unit 0 family inet address 20.20.20.2/30
 set interfaces ge-0/0/2 unit 0 family mpls
 set interfaces lo0 unit 0 family inet address 10.255.10.4/32
 set routing-options router-id 10.255.10.4
 set protocols rsvp interface all
 set protocols rsvp interface fxp0.0 disable
 set protocols mpls interface all
 set protocols mpls interface fxp0.0 disable
 set protocols ospf traffic-engineering
 set protocols ospf area 0.0.0.0 interface all
 set protocols ospf area 0.0.0.0 interface fxp0.0 disable

PE2 set interfaces ge-0/0/0 vlan-tagging
 set interfaces ge-0/0/0 encapsulation flexible-ethernet-services
 set interfaces ge-0/0/0 unit 0 vlan-id 1
 set interfaces ge-0/0/0 unit 0 family inet address 2.2.2.2/30
 set interfaces ge-0/0/0 unit 0 family mpls
 set interfaces ge-0/0/0 unit 10 encapsulation vlan-ccc
 set interfaces ge-0/0/0 unit 10 vlan-id 10
 set interfaces ge-0/0/1 unit 0 family inet address 30.30.30.2/30
 set interfaces ge-0/0/1 unit 0 family mpls
 set interfaces ge-0/0/2 unit 0 family inet address 80.80.80.1/30
 set interfaces ge-0/0/2 unit 0 family mpls
 set interfaces gre unit 0 tunnel source 2.2.2.2
 set interfaces gre unit 0 tunnel destination 2.2.2.1
 set interfaces gre unit 0 family inet address 10.35.101.26/30
 set interfaces lo0 unit 0 family inet address 10.255.10.5/32
 set routing-options router-id 10.255.10.5
 set protocols rsvp associated-bidirectional-lsp single-sided-provisioning
 set protocols rsvp interface all
 set protocols rsvp interface fxp0.0 disable
 set protocols rsvp peer-interface CE2 dynamic-bidirectional-transport
 set protocols mpls interface all
 set protocols mpls interface fxp0.0 disable
 set protocols ospf traffic-engineering
 set protocols ospf area 0.0.0.0 interface all
 set protocols ospf area 0.0.0.0 interface fxp0.0 disable
 set protocols link-management te-link link1 local-address 10.36.1.2
 set protocols link-management te-link link1 remote-address 10.36.1.1
 set protocols link-management te-link link1 ethernet-vlan vlan-id-range 1-1000
 set protocols link-management te-link link1 interface ge-0/0/0
 set protocols link-management peer CE2 address 10.255.10.6
 set protocols link-management peer CE2 control-channel gre.0
 set protocols link-management peer CE2 te-link link1

CE2 set interfaces ge-0/0/0 vlan-tagging
 set interfaces ge-0/0/0 unit 1 vlan-id 1
 set interfaces ge-0/0/0 unit 1 family inet address 2.2.2.1/24

```

```

set interfaces ge-0/0/0 unit 1 family mpls
set interfaces ge-0/0/0 unit 10 vlan-id 10
set interfaces ge-0/0/0 unit 10 family inet address 10.10.10.2/24
set interfaces ge-0/0/0 unit 10 family mpls
set interfaces gre unit 0 tunnel source 2.2.2.1
set interfaces gre unit 0 tunnel destination 2.2.2.2
set interfaces gre unit 0 family inet address 10.35.101.25/30
set interfaces lo0 unit 0 family inet address 10.255.10.6/32
set routing-options router-id 10.255.10.6
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols rsvp peer-interface PE2
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols link-management te-link link10 local-address 10.36.1.1
set protocols link-management te-link link10 remote-address 10.36.1.2
set protocols link-management te-link link10 ethernet-vlan vlan-id-range 1-1000
set protocols link-management te-link link10 interface ge-0/0/0
set protocols link-management peer PE2 address 10.255.10.5
set protocols link-management peer PE2 control-channel gre.0
set protocols link-management peer PE2 te-link link10

```

### Configuring the Client Router

**Step-by-Step Procedure** The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router CE1:



**NOTE:** Repeat this procedure for Router CE2 in the server-layer network, after modifying the appropriate interface names, addresses, and any other parameters for the router.

1. Configure the interface connecting Router CE1 to Router PE1.  

```

[edit interfaces]
user@CE1# set ge-0/0/0 vlan-tagging

```
2. Configure the control VLAN for the ge-0/0/0 interface.  

```

[edit interfaces]
user@CE1# set ge-0/0/0 unit 1 vlan-id 1
user@CE1# set ge-0/0/0 unit 1 family inet address 1.1.1.1/30
user@CE1# set ge-0/0/0 unit 1 family mpls

```
3. Configure the LSP VLAN on the ge-0/0/0 interface.  

```

[edit interfaces]
user@CE1# set ge-0/0/0 unit 10 vlan-id 10
user@CE1# set ge-0/0/0 unit 10 family inet address 10.10.10.1/24
user@CE1# set ge-0/0/0 unit 10 family mpls

```
4. Configure the GRE tunnel as the controlling interface for Router CE1.

- ```
[edit interfaces]
user@CE1# set gre unit 0 tunnel source 1.1.1.1
user@CE1# set gre unit 0 tunnel destination 1.1.1.2
user@CE1# set gre unit 0 family inet address 10.35.100.25/30
```
5. Configure the loopback interface of Router CE1.


```
[edit interfaces]
user@CE1# set lo0 unit 0 family inet address 10.255.10.1/32
```
 6. Configure the loopback address of Router CE1 as its router ID.


```
[edit routing-options]
user@CE1# set router-id 10.255.10.1
```
 7. Enable RSVP on all the interfaces of Router CE1, excluding the management interface.


```
[edit protocols]
user@CE1# set rsvp interface all
user@CE1# set rsvp interface fxp0.0 disable
```
 8. Configure the RSVP peer interface for Router CE1.


```
[edit protocols]
user@CE1# set rsvp peer-interface PE1
```
 9. Disable automatic path computation for label-switched paths (LSPs).


```
[edit protocols]
user@CE1# set mpls no-cspf
```
 10. Configure the LSP to connect Router CE1 to Router CE2.


```
[edit protocols]
user@CE1# set mpls label-switched-path CE1-to-CE2 from 10.255.10.1
user@CE1# set mpls label-switched-path CE1-to-CE2 to 10.255.10.6
```
 11. Configure the CE1-to-CE2 LSP attributes.


```
[edit protocols]
user@CE1# set mpls label-switched-path CE1-to-CE2 lsp-attributes switching-type
    ethernet-vlan
user@CE1# set mpls label-switched-path CE1-to-CE2 lsp-attributes upstream-label
    vlan-id 10
user@CE1# set mpls label-switched-path CE1-to-CE2 bandwidth 100m
```
 12. Configure the CE1-to-CE2 LSP path and path parameters.


```
[edit protocols]
user@CE1# set mpls label-switched-path CE1-to-CE2 primary path1
user@CE1# set mpls path path1 10.35.1.2 strict
user@CE1# set mpls path path1 10.255.10.5 loose
user@CE1# set mpls path path1 10.36.1.1 strict
```
 13. Enable MPLS on all the interfaces of Router CE1, excluding the management interface.


```
[edit protocols]
user@CE1# set mpls interface all
user@CE1# set mpls interface fxp0.0 disable
```

14. Configure a traffic engineering link, and assign addresses for the local and remote end of the link.

```
[edit protocols]
user@CE1# set link-management te-link link10 local-address 10.35.1.1
user@CE1# set link-management te-link link10 remote-address 10.35.1.2
```
15. Enable setting up of Layer 2 VLAN LSP on the link10 traffic engineering link.

```
[edit protocols]
user@CE1# set link-management te-link link10 ethernet-vlan
```
16. Configure the Router CE1 interface as the member interface of the link10 traffic engineering link.

```
[edit protocols]
user@CE1# set link-management te-link link10 interface ge-0/0/0
```
17. Configure Router PE1 as the Link Management Protocol (LMP) peer for Router CE1, and configure the peer attributes.

```
[edit protocols]
user@CE1# set link-management peer PE1 address 10.255.10.2
user@CE1# set link-management peer PE1 control-channel gre.0
user@CE1# set link-management peer PE1 te-link link10
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@CE1# show interfaces
ge-0/0/0 {
  vlan-tagging;
  unit 1 {
    vlan-id 1;
    family inet {
      address 1.1.1.1/30;
    }
    family mpls;
  }
  unit 10 {
    vlan-id 10;
    family inet {
      address 10.10.10.1/24;
    }
    family mpls;
  }
}
gre {
  unit 0 {
    tunnel {
      source 1.1.1.1;
      destination 1.1.1.2;
    }
    family inet {
      address 10.35.100.25/30;
    }
  }
}
```

```

    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.10.1/32;
      }
    }
  }
}

user@CE1# show routing-options
router-id 10.255.10.1;

user@CE1# show protocols
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
    peer-interface PE1;
  }
  mpls {
    no-cspf;
    label-switched-path CE1-to-CE2 {
      from 10.255.10.1;
      to 10.255.10.6;
      lsp-attributes {
        switching-type ethernet-vlan;
        upstream-label {
          vlan-id 10;
        }
      }
    }
    bandwidth 100m;
    primary path1;
  }
  path path1 {
    10.35.1.2 strict;
    10.255.10.5 loose;
    10.36.1.1 strict;
  }
  interface all;
  interface fxp0.0 {
    disable;
  }
}
link-management {
  te-link link10 {
    local-address 10.35.1.1;
    remote-address 10.35.1.2;
    ethernet-vlan;
    interface ge-0/0/0;
  }
  peer PE1 {
    address 10.255.10.2;
    control-channel gre.0;
    te-link link10;
  }
}

```

}

Configuring the Server Router

Step-by-Step Procedure The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Router PE1:



NOTE: Repeat this procedure for Router PE2 in the server-layer network, after modifying the appropriate interface names, addresses, and any other parameters for the router.

1. Configure the interface connecting Router PE1 to Router CE1.

```
[edit interfaces]
user@PE1# set ge-0/0/0 vlan-tagging
user@PE1# set ge-0/0/0 encapsulation flexible-ethernet-services
```
2. Configure the control VLAN for the ge-0/0/0 interface.

```
[edit interfaces]
user@PE1# set ge-0/0/0 unit 1 vlan-id 1
user@PE1# set ge-0/0/0 unit 1 family inet address 1.1.1.2/30
user@PE1# set ge-0/0/0 unit 1 family mpls
```
3. Configure the LSP VLAN on the ge-0/0/0 interface.

```
[edit interfaces]
user@PE1# set ge-0/0/0 unit 10 encapsulation vlan-ccc
user@PE1# set ge-0/0/0 unit 10 vlan-id 10
```
4. Configure the interface connecting Router PE1 to the core routers (Router P1 and Router P2).

```
[edit interfaces]
user@PE1# set ge-0/0/1 unit 0 family inet address 70.70.70.1/30
user@PE1# set ge-0/0/1 unit 0 family mpls

user@PE1# set ge-0/0/2 unit 0 family inet address 20.20.20.1/30
user@PE1# set ge-0/0/2 unit 0 family mpls
```
5. Configure the GRE tunnel as the controlling interface for Router PE1.

```
[edit interfaces]
user@PE1# set gre unit 0 tunnel source 1.1.1.2
user@PE1# set gre unit 0 tunnel destination 1.1.1.1
user@PE1# set gre unit 0 family inet address 10.35.100.26/30
```
6. Configure the loopback interface of Router PE1.

```
[edit interfaces]
user@PE1# set lo0 unit 0 family inet address 10.255.10.2/32
```
7. Configure the loopback address of Router PE1 as its router ID.

- ```
[edit routing-options]
user@PE1# set router-id 10.255.10.2
```
8. Configure an associated bidirectional LSP, and enable unidirectional reverse LSP setup for single-sided provisioned forward LSP.
 

```
[edit protocols]
user@PE1# set rsvp associated-bidirectional-lsp single-sided-provisioning
```
  9. Enable RSVP on all the interfaces of Router PE1, excluding the management interface.
 

```
[edit protocols]
user@PE1# set rsvp interface all
user@PE1# set rsvp interface fxp0.0 disable
```
  10. Configure the RSVP peer interface for Router PE1, and enable dynamic setup of bidirectional packet LSP for transporting nonpacket GMPLS LSP.
 

```
[edit protocols]
user@PE1# set rsvp peer-interface CE1 dynamic-bidirectional-transport
```
  11. Enable MPLS on all the interfaces of Router PE1, excluding the management interface.
 

```
[edit protocols]
user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable
```
  12. Configure OSPF with traffic engineering capabilities.
 

```
[edit protocols]
user@PE1# set ospf traffic-engineering
```
  13. Enable OSPF area 0 on all the interfaces of Router PE1, excluding the management interface.
 

```
[edit protocols]
user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
```
  14. Configure a traffic engineering link, and assign addresses for the local and remote end of the link.
 

```
[edit protocols]
user@PE1# set link-management te-link link1 local-address 10.35.1.2
user@PE1# set link-management te-link link1 remote-address 10.35.1.1
```
  15. Enable setting up of a Layer 2 VLAN LSP for a specific range of VLANs on the link1 traffic engineering link.
 

```
[edit protocols]
user@PE1# set link-management te-link link1 ethernet-vlan vlan-id-range 1-1000
```
  16. Configure the Router PE1 interface as the member interface of the link1 traffic engineering link.
 

```
[edit protocols]
user@CE1# set link-management te-link link1 interface ge-0/0/0
```
  17. Configure Router CE1 as the LMP peer for Router PE1, and configure the peer attributes.

```
[edit protocols]
user@CE1# set link-management peer CE1 address 10.255.10.1
user@CE1# set link-management peer CE1 control-channel gre.0
user@CE1# set link-management peer CE1 te-link link1
```

**Results** From configuration mode, confirm your configuration by entering the **show interfaces**, **show routing-options**, and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@PE1# show interfaces
ge-0/0/0 {
 vlan-tagging;
 encapsulation flexible-ethernet-services;
 unit 1 {
 vlan-id 1;
 family inet {
 address 1.1.1.2/30;
 }
 family mpls;
 }
 unit 10 {
 encapsulation vlan-ccc;
 vlan-id 10;
 }
}
ge-0/0/1 {
 unit 0 {
 family inet {
 address 70.70.70.1/30;
 }
 family mpls;
 }
}
ge-0/0/2 {
 unit 0 {
 family inet {
 address 20.20.20.1/30;
 }
 family mpls;
 }
}
gre {
 unit 0 {
 tunnel {
 source 1.1.1.2;
 destination 1.1.1.1;
 }
 family inet {
 address 10.35.100.26/30;
 }
 }
}
lo0 {
 unit 0 {
 family inet {
```



```

 address 10.255.10.2/32;
 }
}
}
user@PE1# show routing-options
router-id 10.255.10.2;
user@PE1# show protocols
rsvp {
 associated-bidirectional-lsp single-sided-provisioning;
 interface all;
 interface fxp0.0 {
 disable;
 }
 peer-interface CE1 {
 dynamic-bidirectional-transport;
 }
}
mpls {
 interface all;
 interface fxp0.0 {
 disable;
 }
}
ospf {
 traffic-engineering;
 area 0.0.0.0 {
 interface all;
 interface fxp0.0 {
 disable;
 }
 }
}
link-management {
 te-link link1 {
 local-address 10.35.1.2;
 remote-address 10.35.1.1;
 ethernet-vlan {
 vlan-id-range 1-1000;
 }
 interface ge-0/0/0;
 }
 peer CE1 {
 address 10.255.10.1;
 control-channel gre.0;
 te-link link1;
 }
}
}

```

## Verification

Confirm that the configuration is working properly.

- [Verifying the Traffic Engineering Link Status on the Client Routers on page 716](#)
- [Verifying the RSVP Session Status on the Client Routers on page 718](#)

- [Verifying the LSP Status on the Server Router on page 718](#)
- [Verifying the CCC Entries in the MPLS Routing Table of the Server Routers on page 719](#)
- [Verifying End-to-End Connectivity on page 720](#)

#### [Verifying the Traffic Engineering Link Status on the Client Routers](#)

**Purpose** Verify the status of the traffic engineering link configured between Router CE1 and Router CE2.

**Action** From operational mode, run the **show link-management** and the **show link-management te-link detail** commands.

```

user@CE1> show link-management
Peer name: PE1, System identifier: 50740
State: Up, Control address: 10.255.10.2
Hello interval: 150, Hello dead interval: 500
Control-channel State
gre.0 Active
TE links:
link10

TE link name: link10, State: Up
Local identifier: 65075, Remote identifier: 0, Local address: 10.35.1.1, Remote
address: 10.35.1.2, Encoding: Ethernet, Switching: EVPL, Minimum bandwidth: 0bps,

Maximum bandwidth: 1000Mbps, Total bandwidth: 1000Mbps, Available bandwidth:
900Mbps
Name State Local ID Remote ID Bandwidth
Used LSP-name
ge-0/0/0 Up 54183 0 1000Mbps
Yes CE1-to-CE2

user@CE1> show link-management te-link detail
TE link name: link10, State: Up
Local identifier: 65075, Remote identifier: 0, Local address: 10.35.1.1, Remote
address: 10.35.1.2, Encoding: Ethernet, Switching: EVPL, Minimum bandwidth: 0bps,

Maximum bandwidth: 1000Mbps, Total bandwidth: 1000Mbps, Available bandwidth:
900Mbps
Resource: ge-0/0/0, Type: IFD, System identifier: 137, State: Up, Local
identifier: 54183, Remote identifier: 0
Total bandwidth: 1000Mbps, Unallocated bandwidth: 900Mbps
Traffic parameters: Encoding: Ethernet, Switching: EVPL, Granularity: Unknown

Maximum allocations: 4094, Number of allocations: 1, Unique allocations: 1,
In use: Yes
LSP name: CE1-to-CE2, Local label: 10, Remote label: 10, Allocated bandwidth:
100Mbps

user@CE2> show link-management
Peer name: PE2, System identifier: 50743
State: Up, Control address: 10.255.10.5
Hello interval: 150, Hello dead interval: 500
Control-channel State
gre.0 Active
TE links:
link10

TE link name: link10, State: Up
Local identifier: 65075, Remote identifier: 0, Local address: 10.36.1.1, Remote
address: 10.36.1.2, Encoding: Ethernet, Switching: EVPL, Minimum bandwidth: 0bps,

Maximum bandwidth: 1000Mbps, Total bandwidth: 1000Mbps, Available bandwidth:
900Mbps
Name State Local ID Remote ID Bandwidth
Used LSP-name
ge-0/0/0 Up 54183 0 1000Mbps
Yes CE1-to-CE2

```

**Meaning** The Link Management Protocol (LMP) peering has been established between the client routers, and the traffic engineering link is up on both Routers CE1 and CE2.

#### Verifying the RSVP Session Status on the Client Routers

**Purpose** Verify the status of the RSVP sessions between Router CE1 and Router CE2.

**Action** From operational mode, run the **show rsvp session** command.

```
user@CE1> show rsvp session
Ingress RSVP: 1 sessions
 To From State Rt Style Labelin Labelout LSPname
 10.255.10.6 10.255.10.1 Up 0 1 FF - 10 CE1-to-CE2 Bidir
Total 1 displayed, Up 1, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@CE2> show rsvp session
Ingress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress RSVP: 1 sessions
 To From State Rt Style Labelin Labelout LSPname
 10.255.10.6 10.255.10.1 Up 0 1 FF 10 - CE1-to-CE2 Bidir
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

**Meaning** The RSVP sessions are established between the ingress router, Router CE1, and the egress router, Router CE2.

#### Verifying the LSP Status on the Server Router

**Purpose** Verify the status of the MPLS LSP on Router PE1.

**Action** From operational mode, run the **show mpls lsp** command.

```

user@PE1> show mpls lsp
Ingress LSP: 1 sessions
 To From State Rt P ActivePath LSPName
10.255.10.5 10.255.10.2 Up 0 *
vlan:0:10:8176:10.255.10.2->10.255.10.5 Assoc-Bidir
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
 To From State Rt Style Labelin Labelout LSPName
10.255.10.2 10.255.10.5 Up 0 1 FF 3 -
vlan:0:10:8176:10.255.10.2->10.255.10.5:rev
Total 1 displayed, Up 1, Down 0

Transit LSP: 1 sessions
 To From State Rt Style Labelin Labelout LSPName
10.255.10.6 10.255.10.1 Up 0 1 FF 10 299808 CE1-to-CE2 Bidir
Total 1 displayed, Up 1, Down 0

```

**Meaning** The CE1-to-CE2 LSP is established, and the output displays the LSP attributes.

#### Verifying the CCC Entries in the MPLS Routing Table of the Server Routers

**Purpose** Verify the circuit cross-connect (CCC) interface entries in the MPLS routing table.

**Action** From operational mode, run the **show route table mpls.0** and the **show route forwarding-table ccc ccc-interface** commands.

```
user@PE1> show route table mpls.0
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0 *[MPLS/0] 1d 22:14:51, metric 1
 Receive
1 *[MPLS/0] 1d 22:14:51, metric 1
 Receive
2 *[MPLS/0] 1d 22:14:51, metric 1
 Receive
13 *[MPLS/0] 1d 22:14:51, metric 1
 Receive
299824 *[RSVP/7/1] 17:32:07, metric 1
 > via ge-0/0/0.10, Pop
ge-0/0/0.10 *[RSVP/7/1] 17:32:07, metric 1
 > to 20.20.20.2 via ge-0/0/2.0, label-switched-path CE1-to-CE2

user@PE1> show route forwarding-table ccc ge-0/0/0.10
Routing table: default.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
ge-0/0/0.10 (CCC) user 0 20.20.20.2 Push 299808, Push 299872(top)
581 2 ge-0/0/2.0

Routing table: __mpls-oam__.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 534 1
```

**Meaning** The output displays the CCC interface that is the client-router-facing interface and the next-hop details for that interface.

### Verifying End-to-End Connectivity

**Purpose** Verify the connectivity between Router CE1 and the remote client router, Router CE2.

**Action** From operational mode, run the **ping** command.

```
user@CE1> ping 10.10.10.2
PING 10.10.10.2 (10.10.10.2): 56 data bytes
64 bytes from 10.10.10.2: icmp_seq=0 ttl=64 time=15.113 ms
64 bytes from 10.10.10.2: icmp_seq=1 ttl=64 time=13.353 ms
64 bytes from 10.10.10.2: icmp_seq=2 ttl=64 time=13.769 ms
64 bytes from 10.10.10.2: icmp_seq=3 ttl=64 time=10.341 ms
64 bytes from 10.10.10.2: icmp_seq=4 ttl=64 time=12.597 ms
^C
--- 10.10.10.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 10.341/13.035/15.113/1.575 ms
```

**Meaning** The ping from Router CE1 to Router CE2 is successful.

**Related Documentation**

- [GMPLS RSVP-TE VLAN LSP Signaling Overview on page 692](#)





## CHAPTER 19

# Using a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs over a Single RSVP LSP

- [Hierarchy of RSVP LSPs Overview on page 723](#)
- [Hierarchy of RSVP LSPs Terminology on page 723](#)
- [Hierarchy of RSVP LSPs Standard on page 724](#)
- [Hierarchy of RSVP LSPs on page 724](#)
- [Advertising the Forwarding Adjacency with OSPF on page 724](#)
- [Configuring a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs Over a Single RSVP LSP on page 724](#)

## Hierarchy of RSVP LSPs Overview

---

This chapter provides overview information and configuration instructions for hierarchies of RSVP label-switched paths (LSPs), which enable you to tunnel multiple RSVP LSPs over a single RSVP LSP.

The following sections provide an overview of how a hierarchy of RSVP LSPs functions:

- [Hierarchy of RSVP LSPs on page 724](#)
- [Advertising the Forwarding Adjacency with OSPF on page 724](#)

## Hierarchy of RSVP LSPs Terminology

---

### F

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Forwarding adjacency</b>     | A traffic engineering link created by a forwarding adjacency LSP. You can create a forwarding adjacency between two routers in a network by configuring a forwarding adjacency LSP. Forwarding adjacencies can only be statically configured. However, you can configure OSPF to advertise the forwarding adjacency to other routers. When an RSVP LSP traverses a forwarding adjacency, existing MPLS features such as fast reroute continue to function. |
| <b>Forwarding adjacency LSP</b> | An RSVP LSP used to tunnel other RSVP LSPs; forms the basis for a forwarding adjacency.                                                                                                                                                                                                                                                                                                                                                                    |

## Hierarchy of RSVP LSPs Standard

---

For more information on how a hierarchy of RSVP LSPs functions, see RFC 4206, *Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)*.

## Hierarchy of RSVP LSPs

---

Forwarding adjacencies are configured and managed as point-to-point traffic engineering links by including statements at the **[edit protocols link-management]** hierarchy level. For the forwarding adjacency to function properly, you also need to make RSVP aware of the forwarding adjacency by configuring the corresponding peer interface at the **[edit protocols rsvp]** hierarchy level.

Although forwarding adjacency LSPs are configured and managed as traffic engineering links on the local router, it is not necessary to advertise these traffic engineering links to other routers in the network. However, if you want to automatically forward MPLS traffic over the forwarding adjacency or want other routers to compute paths over the forwarding adjacency, you must configure OSPF to advertise the forwarding adjacency to the other routers in the network and add the forwarding adjacency to the traffic engineering database. OSPF is the only supported interior gateway protocol (IGP).

## Advertising the Forwarding Adjacency with OSPF

---

Once a forwarding adjacency LSP and the corresponding traffic engineering link you have configured, you can configure OSPF to advertise the forwarding adjacency. Unlike regular traffic engineering links, OSPF hellos are not exchanged between the forwarding adjacency LSP endpoints and therefore no routing adjacency is created between the forwarding adjacency endpoints. If you issue a **show ospf neighbor** command on an ingress forwarding adjacency, the command displays the egress router of the forwarding adjacency LSP as a neighbor. However, no real OSPF adjacency is established (no OSPF hellos are exchanged) between the ingress and egress routers. For display purposes only, OSPF creates a pseudo-neighbor corresponding to the peer.

You can configure forwarding adjacencies over existing MPLS networks. A forwarding adjacency LSP is signaled as a regular MPLS LSP without generalized MPLS (GMPLS) extensions. When the forwarding adjacency LSP is advertised as a traffic engineering link in OSPF, the corresponding traffic engineering link in OSPF is also advertised as a regular MPLS traffic engineering link without GMPLS extensions.

## Configuring a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs Over a Single RSVP LSP

---

The following sections describe how to configure a hierarchy of RSVP LSPs:

- [Configuring an RSVP LSP on Ingress Routers on page 725](#)
- [Configuring Forwarding Adjacencies on page 725](#)

- [Configuring RSVP for Forwarding Adjacencies on page 726](#)
- [Advertising Forwarding Adjacencies Using OSPF on page 727](#)

## Configuring an RSVP LSP on Ingress Routers

Configure a standard RSVP LSP on the ingress router to be used as the forwarding adjacency LSP (see [label-switched-path](#)). This LSP requires no special configuration to function as a forwarding adjacency LSP.

## Configuring Forwarding Adjacencies

A forwarding adjacency is a type of GMPLS traffic engineering link. It requires that you configure local and remote addresses to identify the link. A forwarding adjacency is associated with a specific peer router. You could configure multiple forwarding adjacencies to the same peer router.

To configure a forwarding adjacency, you need to configure the **te-link** statement at the **[edit protocols link-management]** hierarchy level:

```
[edit protocols link-management]
te-link te-link-name {
 label-switched-path lsp-name;
 local-address ip-address;
 remote-address ip-address;
}
```

For more information on how to configure GMPLS traffic engineering links, see “[Configuring LMP Traffic Engineering Links](#)” on page 678.



**NOTE:** Do not configure the control channel for a forwarding adjacency peer router. Configuring a control channel causes the commit operation to fail.

The following sections describe how to configure the **te-link** statement for a forwarding adjacency:

- [Configuring the Local IP Address for Forwarding Adjacencies on page 725](#)
- [Configuring the Remote IP Address for Forwarding Adjacencies on page 726](#)
- [Configuring the LSP for Forwarding Adjacencies on page 726](#)

### Configuring the Local IP Address for Forwarding Adjacencies

To configure the local IP address for the forwarding adjacency, include the **local-address** statement:

```
local-address ip-address;
```

For a list of hierarchy levels at which you can include this statement, see the statement summary section for this statement.



**NOTE:** If you configure **local-address**, you must also configure **remote-address**.

### Configuring the Remote IP Address for Forwarding Adjacencies

The address of the peer router is the node ID for the forwarding adjacency LSP egress node. You configure this node ID for the forwarding adjacency using the **remote-address** statement:

```
remote-address ip-address;
```

You can include this statement at the following hierarchy levels:

- [edit protocols link-management **te-link** *te-link-name*]
- [edit logical-systems *logical-system-name* protocols link-management **te-link** *te-link-name*]



NOTE: If you configure **remote-address**, you must also configure **local-address**.

### Configuring the LSP for Forwarding Adjacencies

To configure a router to function as a forwarding adjacency, use the **label-switched-path** statement and specify the LSP configured in “Configuring a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs Over a Single RSVP LSP” on page 724:

```
label-switched-path label-switched-path-name;
```

You can include this statement at the following hierarchy levels:

- [edit protocols link-management **te-link** *te-link-name*]
- [edit logical-systems *logical-system-name* protocols link-management **te-link** *te-link-name*]

## Configuring RSVP for Forwarding Adjacencies

For the forwarding adjacency to function properly, RSVP must be made aware of it. Do this by specifying the name of the peer interface corresponding to the link-management peer associated with the forwarding adjacency. Including the **peer-interface** statement at the [edit protocols **rsvp**] hierarchy level enables RSVP to use all of the traffic engineering links configured for that peer. You can also configure RSVP control-plane parameters such as the hello interval and refresh reduction.

To configure RSVP to recognize a forwarding adjacency, include the **peer-interface** statement:

```
peer-interface peer-interface-name {
 disable;
 (aggregate | no-aggregate);
 authentication-key key;
 hello-interval seconds;
 (reliable | no-reliable);
}
```

You can include this statement at the following hierarchy levels:

- **[edit protocols rsvp]**
- **[edit logical-systems *logical-system-name* protocols rsvp]**

For more information on how to configure the **peer-interface** statement, see [“Configuring RSVP and OSPF for LMP Peer Interfaces” on page 685](#).

## Advertising Forwarding Adjacencies Using OSPF

You can allow other routers to dynamically signal paths over a forwarding adjacency LSP by configuring OSPF. This configuration is optional.

If you configure OSPF to advertise a forwarding adjacency LSP, the LSP is added to the traffic engineering database on each router in the traffic engineering domain. Because the forwarding adjacency LSP is unidirectional, the corresponding traffic engineering link (forwarding adjacency) is also unidirectional. The forwarding adjacency LSP appears as a standard traffic engineering database half-link to all routers in the traffic engineering domain.

Constrained Shortest Path First (CSPF) performs a bidirectional link check to ensure that traffic can flow in both directions. CSPF checks for a reverse link, either the exact reverse forwarding adjacency or another reverse link. If there is no reverse link from the forwarding adjacency LSP egress router to the forwarding adjacency LSP ingress router, the CSPF check fails.

CSPF might find another parallel reverse link. However, the LSP cannot function properly over the forwarding adjacency unless you have explicitly configured a corresponding forwarding adjacency LSP to handle the traffic flowing in the opposite direction on the forwarding adjacency LSP egress router.

To advertise the traffic engineering properties of a forwarding adjacency to a specific peer router, include the **peer-interface** statement:

```
peer-interface peer-interface-name {
 dead-interval seconds;
 disable;
 hello-interval seconds;
 retransmit-interval seconds;
 transit-delay seconds;
}
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols ospf area *area-name*]**
- **[edit logical-systems *logical-system-name* protocols ospf area *area-name*]**

For more information on how to configure the **peer-interface** statement, see [“Configuring RSVP and OSPF for LMP Peer Interfaces” on page 685](#).



## PART 7

# Configuring Path Computation Element Protocol (PCEP)

- [PCEP Overview on page 731](#)
- [Configuring PCEP on page 733](#)





# PCEP Overview

- [PCEP Overview on page 731](#)


## PCEP Overview

A Path Computation Element (PCE) is an entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints. A Path Computation Client (PCC) is any client application requesting a path computation to be performed by a PCE. The Path Computation Element Protocol (PCEP) enables communications between a PCC and a PCE, or between two PCEs (defined in RFC 5440).

PCEP is a TCP-based protocol defined by the IETF PCE Working Group, and defines a set of messages and objects used to manage PCEP sessions and to request and send paths for multidomain traffic engineered LSPs (TE LSPs). It provides a mechanism for a PCE to perform path computation for a PCC's external LSPs. The PCEP interactions include LSP status reports sent by the PCC to the PCE, and PCE updates for the external LSPs.

[Figure 77 on page 731](#) illustrates the role of PCEP in the client-side implementation of a stateful PCE architecture in an MPLS RSVP-TE enabled network.

**Figure 77: PCEP Session**



A TCP-based PCEP session connects a PCC to an external PCE. The PCC initiates the PCEP session and stays connected to the PCE for the duration of the PCEP session. During the PCEP session, the PCC requests LSP parameters from the stateful PCE. On

receiving one or more LSP parameters from the PCE, the PCC re-signals the TE LSP. When the PCEP session is terminated, the underlying TCP connection is closed immediately, and the PCC attempts to re-establish the PCEP session.

Thus, the PCEP functions include:

- LSP tunnel state synchronization between a PCC and a stateful PCE—When an active stateful PCE connection is detected, a PCC tries to delegate all LSPs to this PCE in a procedure called LSP state synchronization. PCEP enables synchronization of the PCC LSP state to the PCE.
- Delegation of control over LSP tunnels to a stateful PCE—An active stateful PCE controls one or more LSP attributes for computing paths, such as bandwidth, path (ERO), and priority (setup and hold). PCEP enables such delegation of LSPs for path computation.
- Stateful PCE control of timing and sequence of path computations within and across PCEP sessions—An active stateful PCE modifies one or more LSP attributes, such as bandwidth, path (ERO), and priority (setup and hold). PCEP communicates these new LSP attributes from the PCE to the PCC, after which the PCC re-signals the LSP in the specified path.

**Related  
Documentation**

- [Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE on page 733](#)
- [PCEP Hierarchy Level on page 1125](#)

## CHAPTER 21

# Configuring PCEP

- [Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE on page 733](#)

## Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE

---

- [Support of Path Computation Element Protocol for RSVP-TE Overview on page 733](#)
- [Example: Configuring the Path Computation Element Protocol for MPLS RSVP-TE on page 743](#)

## Support of Path Computation Element Protocol for RSVP-TE Overview

- [Understanding MPLS RSVP-TE on page 733](#)
- [Current MPLS RSVP-TE Limitations on page 735](#)
- [Use of an External Path Computing Entity on page 736](#)
- [Components of External Path Computing on page 737](#)
- [Interaction Between a PCE and a PCC Using PCEP on page 739](#)
- [LSP Behavior with External Computing on page 740](#)
- [Configuration Statements Supported for External Computing on page 742](#)
- [PCE-Controlled LSP Protection on page 742](#)
- [Auto-Bandwidth and PCE-Controlled LSP on page 742](#)
- [Impact of Client-Side PCE Implementation on Network Performance on page 743](#)

## Understanding MPLS RSVP-TE

---

Traffic engineering (TE) deals with performance optimization of operational networks, mainly mapping traffic flows onto an existing physical topology. Traffic engineering provides the ability to move traffic flow away from the shortest path selected by the interior gateway protocol (IGP) and onto a potentially less congested physical path across a network.

For traffic engineering in large dense networks, MPLS capabilities can be implemented, because they potentially provide most of the functionality available from an overlay model, in an integrated manner, and at a lower cost than the currently competing alternatives. The primary reason for implementing MPLS traffic engineering is to control paths along which traffic flows through a network. The main advantage of implementing

MPLS traffic engineering is that it provides a combination of ATM's traffic engineering capabilities along with the class-of-service (CoS) differentiation of IP.

In an MPLS network, data plane information is forwarded using label switching. A packet arriving on a provider edge (PE) router from the customer edge (CE) router is applied labels and forwarded to the egress PE router. The labels are removed at the egress router and forwarded out to the appropriate destination as an IP packet. The label switch routers (LSRs) in the MPLS domain use label distribution protocols to communicate the meaning of labels used to forward traffic between and through the LSRs. RSVP-TE is one such label distribution protocol that enables an LSR peer to learn about the other peer's label mappings.

When both MPLS and RSVP are enabled on a router, MPLS becomes a client of RSVP. The primary purpose of the Junos OS RSVP software is to support dynamic signaling within label-switched paths (LSPs). RSVP reserves resources, such as for IP unicast and multicast flows, and requests quality-of-service (QoS) parameters for applications. The protocol was extended in MPLS traffic engineering to enable RSVP to set up LSPs that can be used for traffic engineering in MPLS networks.

When MPLS and RSVP are combined, labels are associated with RSVP flows. 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 is accomplished using different criteria. The set of packets that are assigned the same label value by a specific node belong to the same forwarding equivalence class (FEC), and effectively define the RSVP flow. When traffic is mapped onto an LSP in this way, the LSP is called an LSP tunnel.

LSP tunnels are a way to establish unidirectional label-switched paths. RSVP-TE builds on the RSVP core protocol by defining new objects and modifying existing objects used in the PATH and RESV objects for LSP establishment. The new objects—LABEL-REQUEST object (LRO), RECORD-ROUTE object (RRO), LABEL object, and EXPLICIT-ROUTE object (ERO)—are optional with respect to the RSVP protocol, except for the LRO and LABEL objects, which are both mandatory for establishing LSP tunnels.

In general, RSVP-TE establishes a label-switched path that ensures frame delivery from ingress to egress router. However, with the new traffic engineering capabilities, the following functions are supported in an MPLS domain:

- Possibility to establish a label-switched path using either a full or partial explicit route (RFC 3209).
- Constraint-based LSP establishment over links that fulfill requirements, such as bandwidth and link properties.
- End-point control, which is associated with establishing and managing LSP tunnels at the ingress and egress routers.
- Link management, which manages link resources to do resource-aware routing of TE LSPs and to program MPLS labels.
- MPLS fast reroute (FRR), which manages the LSPs that need protection and assigns backup tunnel information to these LSPs.

---

### Current MPLS RSVP-TE Limitations


---

Although the RSVP extensions for traffic engineering enable better network utilization and meet requirements of classes of traffic, today's MPLS RSVP-TE protocol suite has several issues inherent to its distributed nature. This causes a number of issues during contention for bisection capacity, especially within an LSP priority class where a subset of LSP shares common setup and hold priority values. The limitations of RSVP-TE include:

- Lack of visibility of individual per-LSP, per-device bandwidth demands—The ingress routers in an MPLS RSVP-TE network establish LSPs without having a global view of the bandwidth demand on the network. Information about network resource utilization is only available as total reserved capacity by traffic class on a per interface basis. Individual LSP state is available locally on each label edge router (LER) for its own LSPs only. As a result, a number of issues related to demand pattern arise, particularly within a common setup and hold priority.
- Asynchronous and independent nature of RSVP signaling—In RSVP-TE, the constraints for path establishment are controlled by an administrator. As such, bandwidth reserved for an LSP tunnel is set by the administrator and does not automatically imply any limit on the traffic sent over the tunnel. Therefore, bandwidth available on a traffic engineering link is the bandwidth configured for the link excluding the sum of all reservations made on the link. Thus, the unsigaled demands on an LSP tunnel lead to service degradation of the LSP requiring excess bandwidth, as well as the other LSPs that comply with the bandwidth requirements of the traffic engineering link.
- LSPs established based on dynamic or explicit path options in the order of preference—The ingress routers in an MPLS RSVP-TE network establish LSPs for demands based on the order of arrival. Because the ingress routers do not have a global view of the bandwidth demand on the network, using the order of preference to establish LSPs can cause traffic to be dropped or LSPs from not being established at all when there is an excess of bandwidth demand.

As an example, [Figure 78 on page 736](#) is configured with MPLS RSVP-TE, in which A and G are the label edge routers (LERs). These ingress routers establish LSPs independently based on the order of demands and have no knowledge or control over each other's LSPs. Routers B, C, and D are intermediate or transit routers that connect to the egress routers E and F.

Figure 78: Example MPLS Traffic Engineering



The ingress routers establish LSPs based on the order in which the demands arrive. If Router G receives two demands of capacity 5 each for G-F, then G signals two LSPs – LSP1 and LSP2 – through G-B-D-F. In the same way, when Router A receives the third demand of capacity 10 for A-E, then it signals an LSP, LSP3, through A-B-C-E. However, if the demand on the A-E LSP increases from 10 to 15, Router A cannot signal LSP3 using the same (A-B-C-E) path, because the B-C link has a lower capacity.

Router A should have signaled the increased demand on LSP3 using the A-B-D-C-E path. Since, LSP1 and LSP2 have utilized the B-D link based on the order of demands received, LSP3 is not signaled.

Thus, although adequate max-flow bandwidth is available for all the LSPs, LSP3 is subject to potentially prolonged service degradation. This is due to Router A's lack of global demand visibility and the lack of systemic coordination in demand placement by the ingress routers A and G.

### Use of an External Path Computing Entity

As a solution to the current limitations found in the MPLS RSVP-TE path computation, an external path computing entity with a global view of per-LSP, per-device demand in the network independent of available capacity is required.

Currently, only online and real-time constraint-based routing path computation is provided in an MPLS RSVP-TE network. Each router performs constraint-based routing calculations independent of the other routers in the network. These calculations are based on currently available topology information—information that is usually recent, but not completely accurate. LSP placements are locally optimized, based on current network status. The MPLS RSVP-TE tunnels are set up using the CLI. An operator configures the TE LSP, which is then signaled by the ingress router.

In addition to the existing traffic engineering capabilities, the MPLS RSVP-TE functionality is extended to include an external path computing entity, called the Path Computation Element (PCE). The PCE computes path for the TE LSPs of ingress routers that have been configured for external control. The ingress router that connects to a PCE is called a Path Computation Client (PCC). The PCC is configured with the Path Computation Client Protocol (PCEP) to facilitate external path computing by a PCE.

For more information, see [“Components of External Path Computing” on page 737](#).

To enable external path computing for a PCC's TE LSPs, include the `lsp-external-controller pccd` statement at the `[edit mpls]` and `[edit mpls lsp lsp-name]` hierarchy levels.

### Components of External Path Computing

The components that make up an external path computing system are:

- [Path Computation Element on page 737](#)
- [Path Computation Client on page 738](#)
- [Path Computation Element Protocol on page 738](#)

#### *Path Computation Element*

A Path Computation Element (PCE) can be any entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints. However, a PCE can compute path for only those TE LSPs of a PCC that have been configured for external control.

A PCE can either be stateful or stateless.

- **Stateful PCE**—A stateful PCE maintains strict synchronization between the PCE and network states (in terms of topology and resource information), along with the set of computed paths and reserved resources in use in the network. In other words, a stateful PCE utilizes information from the traffic engineering database as well as information about existing paths (for example, TE LSPs) in the network when processing new requests from the PCC.

A stateful PCE is of two types:

- **Passive stateful PCE**—Maintains synchronization with the PCC and learns the PCC LSP states to better optimize path calculations, but does not have control over them.
- **Active stateful PCE**—Actively modifies the PCC LSPs, in addition to learning about the PCC LSP states.



**NOTE:** In a redundant configuration with main and backup active stateful PCEs, the backup active stateful PCE cannot modify the attributes of delegated LSPs until it becomes the main PCE at the time of a failover. There is no preempting of PCEs in the case of a switchover. The main PCE is backed by a backup PCE, and when the main PCE goes down, the backup PCE assumes the role of the main PCE and remains as the main PCE even though the previously main PCE comes up again.

A stateful PCE provides the following functions:

- Offers offline LSP path computation.
- Triggers LSP re-route when there is a need to re-optimize the network.
- Changes LSP bandwidth when there is an increase in bandwidth demand from an application.
- Modifies other LSP attributes on the router, such as ERO, setup priority, and hold priority.

A PCE has a global view of the bandwidth demand in the network and maintains a traffic engineered database to perform path computations. It performs statistics collection from all the routers in the MPLS domain using SNMP and NETCONF. This provides a mechanism for offline control of PCC's TE LSPs. Although an offline LSP path computation system can be embedded in a network controller, the PCE acts like a full-fledged network controller that provides control over the PCC's TE LSPs, in addition to computing paths.

Although a stateful PCE allows for optimal path computation and increased path computation success, it requires reliable state synchronization mechanisms, with potentially significant control plane overhead and the maintenance of a large amount of data in terms of states, as in the case of a full mesh of TE LSPs.

- Stateless PCE—A stateless PCE does not remember any computed path, and each set of requests is processed independently of each other (RFC 5440).

#### ***Path Computation Client***

A Path Computation Client (PCC) is any client application requesting a path computation to be performed by a PCE.

A PCC can connect to a maximum of 10 PCEs at one time. The PCC to PCE connection can be a configured static route or a TCP connection that establishes reachability. The PCC assigns each connected PCE a priority number. It sends a message to all the connected PCEs with information about its current LSPs, in a process called LSP state synchronization. For the TE LSPs that have external control enabled, the PCC delegates those LSPs to the main PCE. The PCC elects as the main PCE a PCE with the lowest priority number, or the PCE that it connects to first in the absence of priority number.

The PCC re-signals an LSP based on the computed path it receives from a PCE. When the PCEP session with the main PCE is terminated, the PCC elects a new main PCE, and all delegated LSPs to the previously main PCE are delegated to the newly available main PCE.

#### ***Path Computation Element Protocol***

Path Computation Element Protocol (PCEP) is used for communication between PCC and PCE (as well as between two PCEs) (RFC 5440). PCEP is a TCP-based protocol defined by the IETF PCE Working Group, and defines a set of messages and objects used to manage PCEP sessions and to request and send paths for multidomain TE LSPs. The PCEP interactions include PCC messages as well as notifications of specific states related



to the use of a PCE in the context of MPLS RSVP-TE. When PCEP is used for PCE-to-PCE communication, the requesting PCE assumes the role of a PCC.


Thus, the PCEP functions include:

- LSP tunnel state synchronization between PCC and a stateful PCE.
- Delegation of control over LSP tunnels to a stateful PCE.

### Interaction Between a PCE and a PCC Using PCEP

Figure 79 on page 739 illustrates the relationship between a PCE, PCC, and the role of PCEP in the context of MPLS RSVP-TE.

Figure 79: PCC and RSVP-TE



The PCE to PCC communication is enabled by the TCP-based PCEP. The PCC initiates the PCEP session and stays connected to a PCE for the duration of the PCEP session.

Once the PCEP session is established, the PCC performs the following tasks:

1. LSP state synchronization—The PCC sends information about all the LSPs (local and external) to all connected PCEs. For external LSPs, the PCC sends information about any configuration change, RRO change, state change, and so on, to the PCE.
2. LSP delegation—The PCC then delegates the external LSPs to one PCE, which is the main active stateful PCE. It is only the main PCE that can set parameters for the external LSP. The parameters that the main PCE modifies include bandwidth, path (ERO), and priority (setup and hold). The parameters specified in the local configuration are overridden by the parameters that are set by the main PCE.



**NOTE:** When the PCEP session with the main PCE is terminated, the PCC elects a new main PCE, and all delegated LSPs to the previously main PCE are delegated to the newly available main PCE.


3. LSP signaling—On receiving one or more LSP parameters from the main active stateful PCE, the PCC re-signals the TE LSP based on the PCE provided path. If the PCC fails to set up the LSP, it notifies the PCE of the setup failure and waits for the main PCE to provide new parameters for that LSP, and then re-signals it.

When the PCE specifies a path that is incomplete or has loose hops where only the path end-points are specified, the PCC does not perform local constraint-based routing to find out the complete set of hops. Instead, the PCC provides RSVP with the PCE provided path, as it is, for signaling, and the path gets set up using IGP hop by hop routing.

Considering the topology used in [Figure 78 on page 736](#), [Figure 80 on page 740](#) illustrates the partial client-side PCE implementation in the MPLS RSVP-TE enabled network. The ingress routers A and G are the PCCs that are configured to connect to the external stateful PCE through a TCP connection.

The PCE has a global view of the bandwidth demand in the network and performs external path computations after looking up the traffic engineering database. The active stateful PCE then modifies one or more LSP attributes and sends an update to the PCC. The PCC uses the parameters it receives from the PCE to re-signal the LSP.

**Figure 80: Example PCE for MPLS RSVP-TE**



g040838

This way, the stateful PCE provides a cooperative operation of distributed functionality used to address specific challenges of a shortest interdomain constrained path computation. It eliminates congestion scenarios in which traffic streams are inefficiently mapped onto available resources causing over utilization of some subsets of network resources, while other resources remain under utilized.

### LSP Behavior with External Computing

- [LSP Types on page 741](#)
- [LSP Control Mode on page 741](#)

### ***LSP Types***

In a client-side PCE implementation, there are two types of TE LSPs:

- CLI-controlled LSPs—The LSPs that do not have the **`lsp-external-controller pccd`** statement configured are called CLI-controlled LSPs. Although these LSPs are under local control, the PCC updates the connected PCEs with information about the CLI-controlled LSPs during the initial LSP synchronization process. After the initial LSP synchronization, the PCC informs the PCE of any new and deleted LSPs as well.
- PCE-controlled LSPs—The LSPs that have the **`lsp-external-controller pccd`** statement configured are called PCE-controlled LSPs. The PCC delegates the PCE-controlled LSPs to the main PCE for external path computation.

The PCC informs the PCE about the configured parameters of a PCE-controlled LSP, such as bandwidth, ERO, and priorities. It also informs the PCE about the actual values used for these parameters to set up the LSP including the RRO, when available.

The PCC sends such LSP status reports to the PCE only when a re-configuration has occurred or when there is a change in the ERO, RRO, or status of the PCE-controlled LSPs under external control.

There are two types of parameters that come from the CLI configuration of an LSP for a PCE:

- Parameters that are not overridden by a PCE, and that are applied immediately.
- Parameters that are overridden by a PCE. These parameters include bandwidth, path, and priority (setup and hold values). When the control mode switches from external to local, the CLI configured values for these parameters are applied at the next opportunity to re-signal the LSP. The values are not applied immediately.

The CLI-controlled LSPs coexist with the PCE-controlled LSPs on a PCC.

### ***LSP Control Mode***

In a client-side PCE implementation, there are two types of control modes for a PCC-controlled LSP:

- External—By default, all PCE-controlled LSPs are under external control. When an LSP is under external control, the PCC uses the PCE-provided parameters to set up the LSP.
- Local—A PCE-controlled LSP can come under local control. When the LSP switches from external control to local control, path computation is done using the CLI-configured parameters and constraint-based routing. Such a switchover happens only when there is a trigger to re-signal the LSP. Until then, the PCC uses the PCE-provided parameters to signal the PCE-controlled LSP, although the LSP remains under local control.

A PCE-controlled LSP switches to local control from its default external control mode in cases such as no connectivity to a PCE or when a PCE returns delegation of LSPs back to the PCC.

For more information about CLI-controlled LSPs and PCE-controlled LSPs, see “[LSP Types](#)” on page 741.

### Configuration Statements Supported for External Computing

Table 17 on page 742 lists the MPLS and existing LSP configuration statements that apply to a PCE-controlled LSP.

**Table 17: Applicability of MPLS and Existing LSP Configurations to a PCE-Controlled LSP**

| Support for PCE-Controlled LSP                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Applicable LSP Configuration Statements                                                                                                                                     | Applicable MPLS Configuration Statements                                                                                                                                                        |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| These configuration statements can be configured along with the PCE configuration. However, they take effect only when the local configuration is in use. During PCE control, these configuration statements remain inactive.                                                                                                                                                                                                                                                                                         | <ul style="list-style-type: none"> <li>• admin-group</li> <li>• auto-bandwidth</li> <li>• hop-limit</li> <li>• least-fill</li> <li>• most-fill</li> <li>• random</li> </ul> | <ul style="list-style-type: none"> <li>• admin-group</li> <li>• admin-groups</li> <li>• admin-group-extended</li> <li>• hop-limit</li> <li>• no-cspf</li> <li>• smart-optimize-timer</li> </ul> |
| <p>These configuration statements can be configured along with the PCE configuration, but are overridden by the PCE-controlled LSP attributes. However, when the local configuration is in use, the configured values for these configuration statements are applied.</p> <p><b>NOTE:</b> Changes to the local configuration using the CLI while the LSP is under the control of a stateful PCE does not have any effect on the LSP. These changes come into effect only when the local configuration is applied.</p> | <ul style="list-style-type: none"> <li>• bandwidth</li> <li>• primary</li> <li>• priority</li> </ul>                                                                        | <ul style="list-style-type: none"> <li>• priority</li> </ul>                                                                                                                                    |
| These configuration statements cannot be configured along with the PCE configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                 | <ul style="list-style-type: none"> <li>• p2mp</li> <li>• template</li> </ul>                                                                                                | <ul style="list-style-type: none"> <li>• p2mp-lsp-next-hop</li> </ul>                                                                                                                           |

The rest of the LSP configuration statements are applicable in the same way as for existing LSPs. On configuring any of the above configuration statements for a PCE-controlled LSP, an mpls log message is generated to indicate when the configured parameters take effect.

### PCE-Controlled LSP Protection

The protection paths, including fast reroute and bypass LSPs, are locally computed by the PCC using constraint-based routing. A stateful PCE specifies the primary path (ERO) only. A PCE can also trigger a non-standby secondary path, even if the local configuration does not have a non-standby secondary path for LSP protection.

### Auto-Bandwidth and PCE-Controlled LSP

The auto-bandwidth option does not apply to PCE-controlled LSPs, although LSPs under the control of auto-bandwidth and constraint-based routing can coexist with

PCE-controlled LSPs. The statistics collection for auto-bandwidth takes effect only when the control mode of a PCE-controlled LSP changes from external to local. This can happen in cases such as no connectivity to a PCE or when a PCE returns delegation of LSPs back to the PCC.

### **Impact of Client-Side PCE Implementation on Network Performance**

---

The maintenance of a stateful database can be non-trivial. In a single centralized PCE environment, a stateful PCE simply needs to remember all the TE LSPs that the PCE has computed, the TE LSPs that were actually set up (if this can be known), and when the TE LSPs were torn down. However, these requirements cause substantial control protocol overhead in terms of state, network usage and processing, and optimizing links globally across the network. Thus, the concerns of a stateful PCE implementation include:

- Any reliable synchronization mechanism results in significant control plane overhead. PCEs might synchronize state by communicating with each other, but when TE LSPs are set up using distributed computation performed among several PCEs, the problems of synchronization and race condition avoidance become larger and more complex.
- Out-of-band traffic engineering database synchronization can be complex with multiple PCEs set up in a distributed PCE computation model, and can be prone to race conditions, scalability concerns, and so on.
- Path calculations incorporating total network state is highly complex, even if the PCE has detailed information on all paths, priorities, and layers.

In spite of the above concerns, the partial client-side implementation of the stateful PCE is extremely effective in large traffic engineering systems. It provides rapid convergence and significant benefits in terms of optimal resource usage, by providing the requirement for global visibility of a TE LSP state and for ordered control of path reservations across devices within the system being controlled.

### **Example: Configuring the Path Computation Element Protocol for MPLS RSVP-TE**

This example shows how to enable external path computing by a Path Computation Element (PCE) for traffic engineered label-switched paths (TE LSPs) on a Path Computation Client (PCC). It also shows how to configure the Path Computation Element Protocol (PCEP) on the PCC to enable PCE to PCC communication.

- [Requirements on page 743](#)
- [Overview on page 744](#)
- [Configuration on page 746](#)
- [Verification on page 751](#)

#### **Requirements**

---

This example uses the following hardware and software components:

- Three routers that can be a combination of M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, T Series Core Routers, or PTX Series Transport Router, one of which is configured as a PCC.
- A TCP connection to an external stateful PCE from the PCC.

- Junos OS Release 12.3 or later running on the PCC along with the JSDN add-on package.



**NOTE:** The JSDN add-on package is required to be installed along with the core Junos OS installation package.

Before you begin:

1. Configure the device interfaces.
2. Configure MPLS and RSVP-TE.
3. Configure IS-IS or any other IGP protocol.

---

### Overview

Starting with Junos OS Release 12.3, the MPLS RSVP-TE functionality is extended to provide a partial client-side implementation of the stateful PCE architecture (draft-ietf-pce-stateful-pce) on a PCC.

To enable external path computing by a PCE, include the **lsp-external-controller** statement on the PCC at the **[edit mpls]** and **[edit mpls lsp *lsp-name*]** hierarchy levels.

**lsp-external-controller** *pccd*;

An LSP configured with the **lsp-external-controller** statement is referred to as a PCE-controlled LSP and is under the external control of a PCE by default. An active stateful PCE can override the parameters set from the CLI, such as bandwidth, path (ERO), and priority, for such PCE-controlled LSPs of the PCC.

To enable PCE to PCC communication, configure PCEP on the PCC at the **[edit protocols]** hierarchy level.

**pcep** { ... }

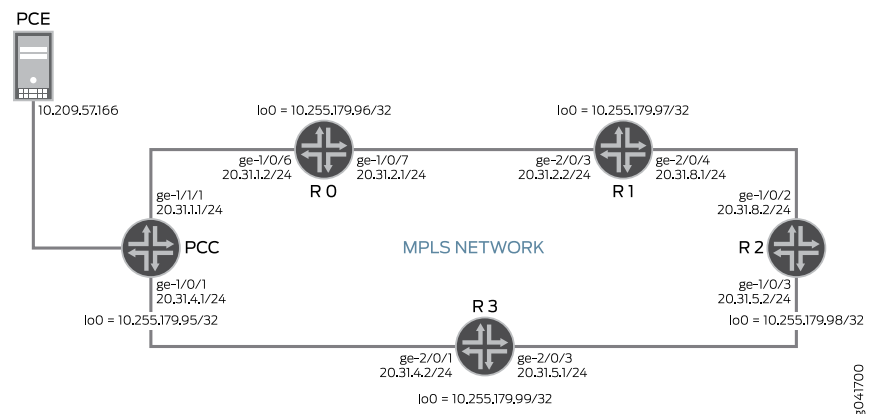
When configuring PCEP on a PCC, be aware of the following considerations:

- The JSDN add-on package is required to be installed along with the core Junos OS installation package.
- Junos OS Release 12.3 supports only stateful PCEs.
- A PCC can connect to a maximum of 10 stateful PCEs. At any given point in time, there can be only one main PCE (the PCE with the lowest priority value, or the PCE that connects to the PCC first in the absence of a PCE priority) to which the PCC delegates LSPs for path computation.
- For Junos OS Release 12.3, the PCC always initiates the PCEP sessions. PCEP sessions initiated by remote PCEs are not accepted by the PCC.
- Existing LSP features, such as LSP protection and make-before-break, work for PCE-controlled LSPs.

- The auto-bandwidth option is turned off for PCE-controlled LSPs, although LSPs under the control of auto-bandwidth and constraint-based routing can coexist with PCE-controlled LSPs.
- PCE-controlled LSPs can be referred to by other CLI configurations, such as `lsp-nexthop` to routes, forwarding adjacencies, CCC connections, and logical tunnels.
- PCE-controlled LSPs do not support GRES.
- PCE-controlled LSPs under logical-systems are not supported.
- PCE-controlled LSPs cannot be point-to-multipoint LSPs.
- Bidirectional LSPs are not supported.
- PCE-controlled LSPs cannot have secondary paths without a primary path.
- PCE-controlled LSPs depend on external path computation, which impacts overall setup time, reroutes, and make-before-break features.
- Setup time and convergence time (reroute, MBB) for existing LSPs is the same as in previous releases, in the absence of PCE-controlled LSPs. However, a small impact is seen in the presence of PCE-controlled LSPs.
- ERO computation time is expected to be significantly higher than local-CSPF.

### Topology

Figure 81: Configuring PCEP for MPLS RSVP-TE



In this example, PCC is the ingress router that connects to the external active stateful PCE.

The external LSPs of Router PCC are computed as follows:

1. Router PCC receives the LSP tunnel configuration that was set up using the CLI. Assuming that the received configuration is enabled with external path computing, Router PCC becomes aware that some of the LSP attributes – bandwidth, path, and priority – are under the control of the stateful PCE and delegates the LSP to the PCE.

In this example, the external LSP is called **PCC-to-R2** and it is being set up from Router PCC to Router R2. The CLI-configured ERO for **PCC-to-R2** is PCC-R0-R1-R2. The bandwidth for **PCC-to-R2** is 10m, and both the setup and hold priority values are 4.

2. Router PCC tries to retrieve the PCE-controlled LSP attributes. To do this, Router PCC sends out a PCRpt message to the stateful PCE stating that the LSP has been configured. The PCRpt message communicates the status of the LSP and contains the local configuration parameters of the LSP.
3. The stateful PCE modifies one or more of the delegated LSP attributes and sends the new LSP parameters to Router PCC through the PCUpd message.
4. On receiving the new LSP parameters, Router PCC sets up a new LSP and re-signals it using the PCE-provided path.

In this example, the PCE-provided ERO for **PCC-to-R2** is PCC-R3-R2. The bandwidth for **PCC-to-R2** is 8m, and both the setup and hold priority values are 3.

5. Router PCC sends a PCRpt with the new RRO to the stateful PCE.

### Configuration

#### CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
PCC set interfaces ge-1/0/1 unit 0 family inet address 20.31.4.1/24
set interfaces ge-1/0/1 unit 0 family iso
set interfaces ge-1/0/1 unit 0 family mpls
set interfaces ge-1/1/1 unit 0 family inet address 20.31.1.1/24
set interfaces ge-1/1/1 unit 0 family iso
set interfaces ge-1/1/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 10.255.179.95/32
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls lsp-external-controller pccd
set protocols mpls label-switched-path PCC-to-R2 to 10.255.179.98
set protocols mpls label-switched-path PCC-to-R2 bandwidth 10m
set protocols mpls label-switched-path PCC-to-R2 priority 4 4
set protocols mpls label-switched-path PCC-to-R2 primary to-R2-path
set protocols mpls label-switched-path PCC-to-R2 lsp-external-controller pccd
set protocols mpls path to-R2-path 20.31.1.2 strict
set protocols mpls path to-R2-path 20.31.2.2 strict
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols isis level 1 disable
set protocols isis interface all
set protocols isis interface fxp0.0 disable
```



```

set protocols isis interface lo0.0
set protocols pcep pce pce1 destination-ipv4-address 10.209.57.166
set protocols pcep pce pce1 destination-port 4189
set protocols pcep pce pce1 pce-type active
set protocols pcep pce pce1 pce-type stateful

R0 set interfaces ge-1/0/6 unit 0 family inet address 20.31.1.2/24
 set interfaces ge-1/0/6 unit 0 family iso
 set interfaces ge-1/0/6 unit 0 family mpls
 set interfaces ge-1/0/7 unit 0 family inet address 20.31.2.1/24
 set interfaces ge-1/0/7 unit 0 family iso
 set interfaces ge-1/0/7 unit 0 family mpls
 set interfaces lo0 unit 0 family inet address 10.255.179.96/32
 set protocols rsvp interface all
 set protocols rsvp interface fxp0.0 disable
 set protocols mpls interface all
 set protocols mpls interface fxp0.0 disable
 set protocols isis level 1 disable
 set protocols isis interface all
 set protocols isis interface fxp0.0 disable
 set protocols isis interface lo0.0

R1 set system ports console log-out-on-disconnect
 set interfaces ge-2/0/3 unit 0 family inet address 20.31.2.2/24
 set interfaces ge-2/0/3 unit 0 family iso
 set interfaces ge-2/0/3 unit 0 family mpls
 set interfaces ge-2/0/4 unit 0 family inet address 20.31.8.1/24
 set interfaces ge-2/0/4 unit 0 family iso
 set interfaces ge-2/0/4 unit 0 family mpls
 set interfaces lo0 unit 0 family inet address 10.255.179.97/32
 set protocols rsvp interface all
 set protocols rsvp interface fxp0.0 disable
 set protocols mpls interface all
 set protocols mpls interface fxp0.0 disable
 set protocols isis level 1 disable
 set protocols isis interface all
 set protocols isis interface fxp0.0 disable
 set protocols isis interface lo0.0

R2 set interfaces ge-1/0/2 unit 0 family inet address 20.31.8.2/24
 set interfaces ge-1/0/2 unit 0 family iso
 set interfaces ge-1/0/2 unit 0 family mpls
 set interfaces ge-1/0/3 unit 0 family inet address 20.31.5.2/24
 set interfaces ge-1/0/3 unit 0 family iso
 set interfaces ge-1/0/3 unit 0 family mpls
 set interfaces lo0 unit 0 family inet address 10.255.179.98/32
 set protocols rsvp interface all
 set protocols rsvp interface fxp0.0 disable
 set protocols mpls interface all
 set protocols mpls interface fxp0.0 disable
 set protocols isis level 1 disable
 set protocols isis interface all
 set protocols isis interface fxp0.0 disable
 set protocols isis interface lo0.0

```

```

R3 set interfaces ge-2/0/1 unit 0 family inet address 20.31.4.2/24
 set interfaces ge-2/0/1 unit 0 family iso
 set interfaces ge-2/0/1 unit 0 family mpls
 set interfaces ge-2/0/3 unit 0 family inet address 20.31.5.1/24
 set interfaces ge-2/0/3 unit 0 family iso
 set interfaces ge-2/0/3 unit 0 family mpls
 set interfaces lo0 unit 0 family inet address 10.255.179.99/32
 set protocols rsvp interface all
 set protocols rsvp interface fxp0.0 disable
 set protocols mpls interface all
 set protocols mpls interface fxp0.0 disable
 set protocols isis level 1 disable
 set protocols isis interface all
 set protocols isis interface fxp0.0 disable
 set protocols isis interface lo0.0

```

**Step-by-Step Procedure** The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure Router PCC:



**NOTE:** Repeat this procedure for every Juniper Networks ingress router in the MPLS domain, after modifying the appropriate interface names, addresses, and any other parameters for each router.

1. Configure the interfaces.

To enable MPLS, include the protocol family on the interface so that the interface does not discard incoming MPLS traffic.

**[edit interfaces]**

```

user@PCC# set ge-1/0/1 unit 0 family inet address 20.31.4.1/24
user@PCC# set ge-1/0/1 unit 0 family iso
user@PCC# set ge-1/0/1 unit 0 family mpls

```

```

user@PCC# set ge-1/1/1 unit 0 family inet address 20.31.1.1/24
user@PCC# set ge-1/1/1 unit 0 family iso
user@PCC# set ge-1/1/1 unit 0 family mpls

```

```

user@PCC# set lo0 unit 0 family inet address 10.255.179.95/32

```

2. Enable RSVP on all interfaces of Router PCC, excluding the management interface.

**[edit protocols]**

```

user@PCC# set rsvp interface all
user@PCC# set rsvp interface fxp0.0 disable

```

3. Configure the label-switched path (LSP) from Router PCC to Router R2 and enable external control of LSPs by the PCE.

**[edit protocols]**

```

user@PCC# set mpls lsp-external-controller pccd

```

```

user@PCC# set mpls label-switched-path PCC-to-R2 to 10.255.179.98/32
user@PCC# set mpls label-switched-path PCC-to-R2 bandwidth 10m
user@PCC# set protocols mpls label-switched-path PCC-to-R2 priority 4 4
user@PCC# set protocols mpls label-switched-path PCC-to-R2 primary to-R2-path
user@PCC# set protocols mpls label-switched-path PCC-to-R2
lsp-external-controller pccd

```

4. Configure the LSP from Router PCC to Router R2, which has local control and is overridden by the PCE-provided LSP parameters.

```

[edit protocols]
user@PCC# set mpls path to-R2-path 20.31.1.2/30 strict
user@PCC# set mpls path to-R2-path 20.31.2.2/30 strict

```

5. Enable MPLS on all interfaces of Router PCC, excluding the management interface.

```

[edit protocols]
user@PCC# set mpls interface all
user@PCC# set mpls interface fxp0.0 disable

```

6. Configure IS-IS on all interfaces of Router PCC, excluding the management interface.

```

[edit protocols]
user@PCC# set isis level 1 disable
user@PCC# set isis interface all
user@PCC# set isis interface fxp0.0 disable
user@PCC# set isis interface lo0.0

```

7. Define the PCE that Router PCC connects to, and configure the IP address of the PCE.

```

[edit protocols]
user@PCC# set pcep pce pce1 destination-ipv4-address 10.209.57.166

```

8. Configure the destination port for Router PCC that connects to a PCE using the TCP-based PCEP.

```

[edit protocols]
user@PCC# set pcep pce pce1 destination-port 4189

```

9. Configure the PCE type.

```

[edit protocols]
user@PCC# set pcep pce pce1 pce-type active
user@PCC# set pcep pce pce1 pce-type stateful

```

### Results

From configuration mode, confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@PCC# show interfaces
ge-1/0/1 {
 unit 0 {
 family inet {
 address 20.31.4.1/24;
 }
 family iso;
 }
}

```

```
 family mpls;
 }
}
ge-1/1/1 {
 unit 0 {
 family inet {
 address 20.31.1.1/24;
 }
 family iso;
 family mpls;
 }
}
lo0 {
 unit 0 {
 family inet {
 address 10.255.179.95/32;
 }
 }
}

user@PCC# show protocols
rsvp {
 interface all;
 interface fxp0.0 {
 disable;
 }
}
mpls {
 lsp-external-controller pccd;
 label-switched-path PCC-to-R2 {
 to 10.255.179.98;
 bandwidth 10m;
 priority 4 4;
 primary to-R2-path;
 lsp-external-controller pccd;
 }
 path to-R2-path {
 20.31.1.2 strict;
 20.31.2.2 strict;
 }
 interface all;
 interface fxp0.0 {
 disable;
 }
}
isis {
 level 1 disable;
 interface all;
 interface fxp0.0 {
 disable;
 }
 interface lo0.0;
}
pcep {
 pce pce1 {
 destination-ipv4-address 10.209.57.166;
```

```

 destination-port 4189;
 pce-type active stateful;
 }
}

```

If you are done configuring the device, enter **commit** from configuration mode.

### Verification

Confirm that the configuration is working properly.

- [Verifying the PCEP Session Status on page 751](#)
- [Verifying the PCE-Controlled LSP Status When LSP Control Is External on page 752](#)
- [Verifying the PCE-Controlled LSP Status When LSP Control Is Local on page 753](#)

#### *Verifying the PCEP Session Status*

**Purpose** Verify the PCEP session status between the PCE and Router PCC when the PCE status is up.

**Action** From operational mode, run the **show path-computation-client active-pce** command.

```
user@PCC> show path-computation-client active-pce
```

```
PCE pce1
```

```
General
```

```

IP address : 10.209.57.166
Priority : 0
PCE status : PCE_STATE_UP
Session type : PCE_TYPE_STATEFULACTIVE
PCE-mastership : main

```

```
Counters
```

|           |          |              |              |
|-----------|----------|--------------|--------------|
| PCReqs    | Total: 0 | last 5min: 0 | last hour: 0 |
| PCReps    | Total: 0 | last 5min: 0 | last hour: 0 |
| PCRpts    | Total: 5 | last 5min: 5 | last hour: 5 |
| PCUpdates | Total: 1 | last 5min: 1 | last hour: 1 |

```
Timers
```

|        |                  |        |             |         |
|--------|------------------|--------|-------------|---------|
| Local  | Keepalive timer: | 30 [s] | Dead timer: | 120 [s] |
| Remote | Keepalive timer: | 30 [s] | Dead timer: | 120 [s] |

```
Errors
```

```

PCErr-recv
PCErr-sent
PCE-PCC-NTFS
PCC-PCE-NTFS

```

**Meaning** The output displays information about the current active stateful PCE that Router PCC is connected to. The **PCE status** output field indicates the current status of the PCEP session between the PCE and Router PCC.

For **pce1**, the status of the PCEP session is **PCE\_STATE\_UP**, which indicates that the PCEP session has been established between the PCEP peers.

The statistics of **PCRpts** indicate the number of messages sent by Router PCC to the PCE to report the current status of LSPs. The **PCUpdates** statistics indicate the number of messages received by Router PCC from the PCE. The **PCUpdates** messages include the PCE modified parameters for the PCE-controlled LSPs.

### *Verifying the PCE-Controlled LSP Status When LSP Control Is External*

**Purpose** Verify the status of the PCE-controlled LSP from Router PCC to Router R2 when the LSP is under external control.

**Action** From operational mode, run the **show mpls lsp name PCC-to-R2 extensive** command.

```
user@PCC> show mpls lsp name PCC-to-R2 extensive
Ingress LSP: 1 sessions
```

```
10.255.179.98
 From: 10.255.183.59, State: Up, ActiveRoute: 0, LSPname: PCC-to-R2
 ActivePath: to-R2-path (primary)
 LSPtype: Externally controlled, Penultimate hop popping
 LSP Control Status: Externally controlled
 LoadBalance: Random
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary to-R2-path State: Up
 Priorities: 3 3
 Bandwidth: 8Mbps
 SmartOptimizeTimer: 180
 No computed ERO.
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
 20.31.4.2 20.31.5.2
 21 Mar 11 05:00:56.736 EXTCTRL LSP: Sent Path computation request and LSP
 status
 20 Mar 11 05:00:56.736 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 19 Mar 11 05:00:56.735 Selected as active path
 18 Mar 11 05:00:56.734 EXTCTRL LSP: Sent Path computation request and LSP
 status
 17 Mar 11 05:00:56.734 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 16 Mar 11 05:00:56.734 Record Route: 20.31.4.2 20.31.5.2
 15 Mar 11 05:00:56.734 Up
 14 Mar 11 05:00:56.713 EXTCTRL LSP: Sent Path computation request and LSP
 status
 13 Mar 11 05:00:56.713 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 12 Mar 11 05:00:56.712 Originate Call
 11 Mar 11 05:00:56.712 EXTCTRL_LSP: Received setup parameters : 20.31.4.2
 20.31.5.2
 10 Mar 11 05:00:49.283 EXTCTRL LSP: Sent Path computation request and LSP
 status
 9 Mar 11 05:00:49.283 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 8 Mar 11 05:00:20.581 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 7 Mar 11 05:00:20.581 EXTCTRL LSP: Sent Path computation request and LSP
```

```

status
 6 Mar 11 05:00:20.581 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 5 Mar 11 05:00:20.580 EXTCTRL_LSP: Control status became external
 4 Mar 11 05:00:03.716 EXTCTRL_LSP: Control status became local
 3 Mar 11 05:00:03.714 EXTCTRL_LSP: Sent Path computation request and LSP
status
 2 Mar 11 05:00:03.714 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 1 Mar 11 05:00:00.279 EXTCTRL_LSP: Awaiting external controller connection
 Created: Mon Mar 11 05:00:00 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

**Meaning** In the output, the **LSptype** and **LSP Control Status** output fields show that the LSP is externally controlled. The output also shows a log of the PCEP messages sent between Router PCC and the PCE.

The PCEP session between the PCE and Router PCC is up, and Router PCC receives the following PCE-controlled LSP parameters:

- ERO (path)—20.31.4.2 and 20.31.5.2
- Bandwidth—8Mbps
- Priorities—3 3 (setup and hold values)

#### *Verifying the PCE-Controlled LSP Status When LSP Control Is Local*

**Purpose** Verify the status of the PCE-controlled LSP from Router PCC to Router R2 when the LSP control becomes local.

**Action** From operational mode, run the **show mpls lsp name PCC-to-R2 extensive** command.

```

user@PCC> show mpls lsp name PCC-to-R2 extensive
Ingress LSP: 1 sessions

10.255.179.98
 From: 10.255.183.59, State: Up, ActiveRoute: 0, LSPname: PCC-to-R2
 ActivePath: to-R2-path (primary)
 LSptype: Externally controlled, Penultimate hop popping
 LSP Control Status: Locally controlled
 LoadBalance: Random
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary to-R2-path State: Up
 Priorities: 4 4 (ActualPriorities 3 3)
 Bandwidth: 10Mbps (ActualBandwidth: 8Mbps)
 SmartOptimizeTimer: 180
 No computed ERO.
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
 20.31.4.2 20.31.5.2
 22 Mar 11 05:02:09.618 EXTCTRL_LSP: Control status became local
 21 Mar 11 05:00:56.736 EXTCTRL_LSP: Sent Path computation request and LSP

```

```

status
 20 Mar 11 05:00:56.736 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 19 Mar 11 05:00:56.735 Selected as active path
 18 Mar 11 05:00:56.734 EXTCTRL LSP: Sent Path computation request and LSP
status
 17 Mar 11 05:00:56.734 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 16 Mar 11 05:00:56.734 Record Route: 20.31.4.2 20.31.5.2
 15 Mar 11 05:00:56.734 Up
 14 Mar 11 05:00:56.713 EXTCTRL LSP: Sent Path computation request and LSP
status
 13 Mar 11 05:00:56.713 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 12 Mar 11 05:00:56.712 Originate Call
 11 Mar 11 05:00:56.712 EXTCTRL_LSP: Received setup parameters : 20.31.4.2
20.31.5.2
 10 Mar 11 05:00:49.283 EXTCTRL LSP: Sent Path computation request and LSP
status
 9 Mar 11 05:00:49.283 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 8 Mar 11 05:00:20.581 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 7 Mar 11 05:00:20.581 EXTCTRL LSP: Sent Path computation request and LSP
status
 6 Mar 11 05:00:20.581 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 5 Mar 11 05:00:20.580 EXTCTRL_LSP: Control status became external
 4 Mar 11 05:00:03.716 EXTCTRL_LSP: Control status became local
 3 Mar 11 05:00:03.714 EXTCTRL LSP: Sent Path computation request and LSP
status
 2 Mar 11 05:00:03.714 EXTCTRL_LSP: Computation request/lsp status contains:
 bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 1 Mar 11 05:00:00.279 EXTCTRL LSP: Awaiting external controller connection
 Created: Mon Mar 11 05:00:00 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

**Meaning** In the output, the **LSP Control Status** output field shows that the LSP is under local control. Although the PCE-controlled LSP is under local control, Router PCC continues to use the PCE-provided parameters, until the next opportunity to re-signal the LSP.

The output now displays the LSP parameters that were configured using the CLI along with the PCE-provided parameters used to establish the LSP as the actual values in use.

- Bandwidth—10Mbps (ActualBandwidth: 8Mbps)
- Priorities—4 4 (ActualPriorities 3 3)

On the trigger to re-signal the LSP, Router PCC uses the local configuration parameters to establish the PCE-controlled LSP.

```

user@PCC> show mpls lsp name PCC-to-R2 extensive externally-controlled
Ingress LSP: 1 sessions

```



```

10.255.179.98
 From: 10.255.183.59, State: Up, ActiveRoute: 0, LSPname: PCC-to-R2
 ActivePath: to-R2-path (primary)
 LSPTYPE: Externally controlled, Penultimate hop popping
 LSP Control Status: Locally controlled
 LoadBalance: Random
 Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary to-R2-path State: Up
 Priorities: 4 4
 Bandwidth: 10Mbps
 SmartOptimizeTimer: 180
 Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
20.31.1.2 S 20.31.2.2 S 20.31.8.2 S
 Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
 20.31.1.2 20.31.2.2 20.31.8.2
 28 Mar 11 05:02:51.787 Record Route: 20.31.1.2 20.31.2.2 20.31.8.2
 27 Mar 11 05:02:51.787 Up
 26 Mar 11 05:02:51.697 EXTCTRL_LSP: Applying local parameters with this
signalling attempt
 25 Mar 11 05:02:51.697 Originate Call
 24 Mar 11 05:02:51.696 Clear Call
 23 Mar 11 05:02:51.696 CSPF: computation result accepted 20.31.1.2 20.31.2.2
20.31.8.2
 22 Mar 11 05:02:09.618 EXTCTRL_LSP: Control status became local
 21 Mar 11 05:00:56.736 EXTCTRL LSP: Sent Path computation request and LSP
status
 20 Mar 11 05:00:56.736 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 19 Mar 11 05:00:56.735 Selected as active path
 18 Mar 11 05:00:56.734 EXTCTRL LSP: Sent Path computation request and LSP
status
 17 Mar 11 05:00:56.734 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 16 Mar 11 05:00:56.734 Record Route: 20.31.4.2 20.31.5.2
 15 Mar 11 05:00:56.734 Up
 14 Mar 11 05:00:56.713 EXTCTRL LSP: Sent Path computation request and LSP
status
 13 Mar 11 05:00:56.713 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 12 Mar 11 05:00:56.712 Originate Call
 11 Mar 11 05:00:56.712 EXTCTRL_LSP: Received setup parameters : 20.31.4.2
20.31.5.2
 10 Mar 11 05:00:49.283 EXTCTRL LSP: Sent Path computation request and LSP
status
 9 Mar 11 05:00:49.283 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 8 Mar 11 05:00:20.581 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 7 Mar 11 05:00:20.581 EXTCTRL LSP: Sent Path computation request and LSP
status
 6 Mar 11 05:00:20.581 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 5 Mar 11 05:00:20.580 EXTCTRL_LSP: Control status became external
 4 Mar 11 05:00:03.716 EXTCTRL_LSP: Control status became local
 3 Mar 11 05:00:03.714 EXTCTRL LSP: Sent Path computation request and LSP
status
 2 Mar 11 05:00:03.714 EXTCTRL_LSP: Computation request/lsp status contains:
bandwidth 10000000 priority - setup 4 hold 4 hops: 20.31.1.2 20.31.2.2
 1 Mar 11 05:00:00.279 EXTCTRL LSP: Awaiting external controller connection
Created: Mon Mar 11 05:00:00 2013

```

Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

The **Computed ERO** is 20.31.1.2, 20.31.2.2, and 20.31.8.2. The PCE-controlled LSP is established using the local configuration parameters.

## PART 8

# Troubleshooting Information

- [Troubleshooting MPLS on page 759](#)



## CHAPTER 22

# Troubleshooting MPLS

- [Verify MPLS Interfaces on page 759](#)
- [Verify That Node-Link Protection Is Up on page 761](#)
- [Verify That Link Protection Is Up on page 768](#)
- [Verify One-to-One Backup on page 772](#)
- [Verify That the Primary Path Is Operational on page 779](#)
- [Verify That the Secondary Path Is Established on page 780](#)
- [Checklist for Checking the MPLS Layer on page 782](#)
- [Checking the MPLS Layer on page 783](#)
- [Checklist for Working with the Layered MPLS Troubleshooting Model on page 798](#)
- [Understanding the Layered MPLS Troubleshooting Model on page 798](#)
- [Verify That Load Balancing Is Working on page 805](#)
- [Verify the Operation of Uneven Bandwidth Load Balancing on page 808](#)

### Verify MPLS Interfaces

---

- Purpose** If the MPLS protocol is not configured correctly on the routers in your network, the interfaces are not able to perform MPLS switching.
- Action** To verify MPLS interfaces, enter the following Junos OS command-line interface (CLI) operational mode command:
- ```
user@host> show mpls interface
```

Sample Output 1

The following sample output is for all routers in the network shown in *MPLS Network Topology*.

```
user@R1> show mpls interface
Interface      State      Administrative groups
so-0/0/0.0     Up         <none>
so-0/0/1.0     Up         <none>
so-0/0/2.0     Up         <none>

user@R2> show mpls interface
Interface      State      Administrative groups
so-0/0/0.0     Up         <none>
```

```

so-0/0/1.0      Up      <none>
so-0/0/2.0      Up      <none>
so-0/0/3.0      Up      <none>

```

```
user@R3> show mpls interface
```

```

Interface      State      Administrative groups
so-0/0/0.0      Up          <none>
so-0/0/1.0      Up          <none>
so-0/0/2.0      Up          <none>
so-0/0/3.0      Up          <none>

```

```
user@R4> show mpls interface
```

```

Interface      State      Administrative groups
so-0/0/0.0      Up          <none>
so-0/0/1.0      Up          <none>
so-0/0/2.0      Up          <none>
so-0/0/3.0      Up          <none>

```

```
user@R5> show mpls interface
```

```

Interface      State      Administrative groups
so-0/0/0.0      Up          <none>
so-0/0/1.0      Up          <none>
so-0/0/2.0      Up          <none>

```

```
user@R6> show mpls interface
```

```

Interface      State      Administrative groups
so-0/0/0.0      Up          <none>
so-0/0/1.0      Up          <none>
so-0/0/2.0      Up          <none>
so-0/0/3.0      Up          <none>

```

Sample Output 2

```
user@R6> show mpls interface
```

```

Interface      State      Administrative groups
so-0/0/0.0      Up          <none>
so-0/0/1.0      Up          <none>
so-0/0/3.0      Up          <none>      # so-0/0/2.0 is missing

```

Sample Output 3

```
user@host> show mpls interface
```

```
MPLS not configured
```

Meaning Sample Output 1 shows that all MPLS interfaces on all routers in the network are enabled (**Up**) and can perform MPLS switching. If you fail to configure the correct interface at the **[edit protocols mpls]** hierarchy level or include the **family mpls** statement at the **[edit interfaces type-fpc/pic/port unit number]** hierarchy level, the interface cannot perform MPLS switching, and does not appear in the output for the **show mpls interface** command.

Administrative groups are not configured on any of the interfaces shown in the example network in *MPLS Network Topology*. However, if they were, the output would indicate which affinity class bits are enabled on the router.

Sample Output 2 shows that interface **so-0/0/2.0** is missing and therefore might be incorrectly configured. For example, the interface might not be included at the **[edit**

`protocols mpls`] hierarchy level, or the `family mpls` statement might not be included at the `[edit interfaces type-fpc/pic/port unit number]` hierarchy level. If the interface is configured correctly, RSVP might not have signaled over this interface yet. For more information on determining which interface is incorrectly configured, see *Verify Protocol Families*.

Sample Output 3 shows that the MPLS protocol is not configured at the `[edit protocols mpls]` hierarchy level.

For more information on configuring MPLS on routers in your network, see *Configuring MPLS on Your Network*

Verify That Node-Link Protection Is Up

Purpose After you configure node-link protection, you must check that bypass paths are up. You can also check the number of LSPs protected by the bypass paths. In the network shown in *Node-Link Protection*, two bypass paths should be up: one next-hop bypass path protecting the link between R1 and R2 (or next-hop 10.0.12.14), and a next-next-hop bypass path avoiding R2.

Action To verify node-link protection (many-to-one backup), enter the following Junos OS CLI operational mode commands on the ingress router. You can also issue the commands on transit routers and other routers used in the bypass path for slightly different information.

```
show mpls lsp (See Sample Output on page ?)
show mpls lsp extensive (See Sample Output on page 763)
show rsvp interface (See Sample Output on page 764)
show rsvp interface extensive (See Sample Output on page 765)
show rsvp session detail (See Sample Output on page 766)
```

Sample Output

```
user@R1> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPName
192.168.5.1  192.168.1.1  Up    0 via-r2          *      lsp2-r1-to-r5
Total 1 displayed, Up 1 , Down 0

Egress LSP: 1 sessions
To          From          State Rt Style Labelin Labelout LSPName
192.168.1.1  192.168.5.1  Up    0  1 FF      3      - r5-to-r1
Total 1 displayed, Up 1 , Down 0

Transit LSP: 2 sessions
To          From          State Rt Style Labelin Labelout LSPName
192.168.0.1  192.168.6.1  Up    0  1 FF 100464 101952 lsp1-r6-to-r0
192.168.6.1  192.168.0.1  Up    0  1 FF 100448      3  r0-to-t6
Total 2 displayed, Up 2, Down 0
```

Meaning Sample output from R1 for the `show mpls lsp` command shows a brief description of the state of configured and active LSPs for which R1 is the ingress, transit, and egress router. All LSPs are up. R1 is the ingress router for `lsp2-r1-to-r5`, and the egress router for return LSP `r5-to-r1`. Two LSPs transit R1, `lsp1-r6-to-r0` and the return LSP `r0-to-t6`. For more

detailed information about the LSP, include the **extensive** option when you issue the **show mpls lsp** command.


```

Sample Output user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up , ActiveRoute: 0, LSPname: lsp2-r1-to-r5
  ActivePath: via-r2 (primary)
  Node/Link protection desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary via-r2 State: Up
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
  10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.0.12.14(Label=101872) 10.0.24.2(Label=101360) 10.0.45.2(Label=3)
    11 Jul 11 14:30:58 Link-protection Up
    10 Jul 11 14:28:28 Selected as active path
    [...Output truncated...]
    Created: Tue Jul 11 14:22:58 2006
  Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.1.1
  From: 192.168.5.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r5-to-r1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 146, Since: Tue Jul 11 14:28:36 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 29228 protocol 0
  PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 362 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.45.2 10.0.24.2 10.0.12.14 <self>
  Total 1 displayed, Up 1, Down 0

Transit LSP: 2 sessions

192.168.0.1
  From: 192.168.6.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp1-r6-to-r0, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101952
  Resv style: 1 SE, Label in: 100464, Label out: 101952
  Time left: 157, Since: Tue Jul 11 14:31:38 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 11131 protocol 0
  Node/Link protection desired
  Type: Node/Link protected LSP, using Bypass->10.0.12.14->10.0.24.2
    1 Jul 11 14:31:38 Node protection up, using Bypass->10.0.12.14->10.0.24.2
  PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 509 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 356 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 358 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
  Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

```

```

192.168.6.1
  From: 192.168.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r0-to-t6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100448, Label out: 3
  Time left: 147, Since: Tue Jul 11 14:31:36 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 23481 protocol 0
  PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 358 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.16.2 (so-0/0/3.0) 350 pkts
  RESV rcvfrom: 10.0.16.2 (so-0/0/3.0) 323 pkts
  Explct route: 10.0.16.2
  Record route: 10.0.50.2 10.0.45.2 10.0.24.2 10.0.12.14 <self> 10.0.16.2
Total 2 displayed, Up 2, Down 0

```

Meaning Sample output from R1 for the **show mpls lsp extensive** command shows detailed information about all LSPs for which R1 is the ingress, egress, or transit router, including all past state history and the reason why an LSP failed. All LSPs are up. The main two LSPs **lsp2-r1-to-r5** and **lsp1-r6-to-r0** have node-link protection as indicated by the **Node/Link protection desired** field in the ingress and transit sections of the output. In the ingress section of the output, the **Link-protection Up** field shows that **lsp2-r1-to-r5** has link protection up. In the transit section of the output, the **Type: Node/Link protected LSP** field shows that **lsp1-r6-to-r0** has node-link protection up, and in case of failure will use the bypass **LSP Bypass->10.0.12.14->10.0.24.2**.

Sample Output

```

user@R1> show rsvp interface
RSVP interface: 4 active

```

Interface	State	Active resv	Subscr-ption	Static BW	Available BW	Reserved BW	Highwater mark
fe-0/1/0.0	Up	2	100%	100Mbps	100Mbps	0bps	0bps
fe-0/1/1.0	Up	1	100%	100Mbps	100Mbps	0bps	0bps
fe-0/1/2.0	Up	0	100%	100Mbps	100Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

Meaning Sample output from R1 for the **show rsvp interface** command shows four interfaces enabled with RSVP (**Up**). Interface **fe-0/1/0.0** has two active RSVP reservations (**Active resv**) that might indicate sessions for the two main LSPs, **lsp1-r6-to-r0** and **lsp2-r1-to-r5**. Interface **fe-0/1/0.0** is the connecting interface between R1 and R2, and both LSPs are configured with a strict path through **fe-0/1/0.0**. For more detailed information about what is happening on interface **fe-0/1/0.0**, issue the **show rsvp interface extensive** command.

Sample Output

```

user@R1> show rsvp interface extensive
RSVP interface: 3 active
fe-0/1/0.0 Index 67, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, LinkProtection
  HelloInterval 9(second)
  Address 10.0.12.13
  ActiveResv 2, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = ct0, StaticBW 100Mbps
  ct0: StaticBW 100Mbps, AvailableBW 100Mbps
    MaxAvailableBW 100Mbps = (bc0*subscription)
    ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Protection: On, Bypass: 2, LSP: 2, Protected LSP: 2, Unprotected LSP: 0
    2 Jul 14 14:49:40 New bypass Bypass->10.0.12.14
    1 Jul 14 14:49:34 New bypass Bypass->10.0.12.14->10.0.24.2
  Bypass: Bypass->10.0.12.14, State: Up, Type: LP, LSP: 0, Backup: 0
    3 Jul 14 14:49:42 Record Route: 10.0.17.14 10.0.27.1
    2 Jul 14 14:49:42 Up
    1 Jul 14 14:49:42 CSPF: computation result accepted
  Bypass: Bypass->10.0.12.14->10.0.24.2, State: Up, Type: NP, LSP: 2, Backup: 0
    4 Jul 14 14:50:04 Record Route: 10.0.17.14 10.0.79.2 10.0.59.1 10.0.45.1
    3 Jul 14 14:50:04 Up
    2 Jul 14 14:50:04 CSPF: computation result accepted
    1 Jul 14 14:49:34 CSPF failed: no route toward 10.0.24.2
[...Output truncated...]

```

Meaning Sample output from R1 for the **show rsvp interface extensive** command shows more detailed information about the activity on all RSVP interfaces (3). However, only output for **fe-0/1/0.0** is shown. Protection is enabled (**Protection: On**), with two bypass paths (**Bypass: 2**) protecting two LSPs (**Protected LSP: 2**). All LSPs are protected, as indicated by the **Unprotected LSP: 0** field. The first bypass **Bypass->10.0.12.14** is a link protection bypass path (**Type: LP**), protecting the link between R1 and R2 **fe-0/1/0.0**. The second bypass path **10.0.12.14->10.0.24.2** is a node-link protected LSP, avoiding R2 in case of node failure.

Sample Output user@R1> show rsvp session detail
Ingress RSVP: 2 sessions

```
192.168.4.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->10.0.12.14->10.0.24.2
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 102000
  Resv style: 1 SE, Label in: -, Label out: 102000
  Time left: -, Since: Tue Jul 11 14:30:53 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 60120 protocol 0
  Type: Bypass LSP
  Number of data route tunnel through: 2
  Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.17.14 (fe-0/1/1.0) 336 pkts
  RESV rcvfrom: 10.0.17.14 (fe-0/1/1.0) 310 pkts
  Explct route: 10.0.17.14 10.0.79.2 10.0.59.1 10.0.45.1
  Record route: <self> 10.0.17.14 10.0.79.2 10.0.59.1 10.0.45.1
```

```
192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp2-r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101872
  Resv style: 1 SE, Label in: -, Label out: 101872
  Time left: -, Since: Tue Jul 11 14:28:28 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 60118 protocol 0
  Node/Link protection desired
  Type: Node/Link protected LSP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 344 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 349 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.45.2
  Record route: <self> 10.0.12.14 10.0.24.2 10.0.45.2
Total 2 displayed, Up 2, Down 0
```

Egress RSVP: 1 sessions

```
192.168.1.1
  From: 192.168.5.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r5-to-r1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 147, Since: Tue Jul 11 14:28:36 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 29228 protocol 0
  PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 348 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.45.2 10.0.24.2 10.0.12.14 <self>
Total 1 displayed, Up 1, Down 0
```

Transit RSVP: 2 sessions

```

192.168.0.1
  From: 192.168.6.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp1-r6-to-r0, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101952
  Resv style: 1 SE, Label in: 100464, Label out: 101952
  Time left: 134, Since: Tue Jul 11 14:31:38 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 11131 protocol 0
  Node/Link protection desired
  Type: Node/Link protected LSP
  PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 488 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 339 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 343 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
  Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

192.168.6.1
  From: 192.168.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r0-to-t6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100448, Label out: 3
  Time left: 158, Since: Tue Jul 11 14:31:36 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 23481 protocol 0
  PATH rcvfrom: 10.0.12.14 (fe-0/1/0.0) 344 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.16.2 (so-0/0/3.0) 337 pkts
  RESV rcvfrom: 10.0.16.2 (so-0/0/3.0) 310 pkts
  Explct route: 10.0.16.2
  Record route: 10.0.50.2 10.0.45.2 10.0.24.2 10.0.12.14 <self> 10.0.16.2
Total 2 displayed, Up 2, Down 0

```

Meaning Sample output from **R1** shows detailed information about the RSVP sessions active on **R1**. All sessions are up, with two ingress sessions, one egress session, and two transit sessions.

Within the ingress section, the first session is a bypass path, as indicated by the **Type: Bypass LSP** field; and the second session is a protected LSP (**lsp2-r1-to-r5**) originating on **R1**, as indicated by the **Type: Node/Link protected LSP** field.

Conclusion Multiprotocol Label Switching (MPLS) label-switched path (LSP) link protection and node-link protection are facility-based methods used to reduce the amount of time needed to reroute LSP traffic. These protection methods are often compared to fast reroute—the other Junos OS LSP protection method.

While fast reroute protects LSPs on a one-to-one basis, link protection and node-link protection protect multiple LSPs by using a single, logical bypass LSP. Link protection provides robust backup support for a link, node-link protection bypasses a node or a link, and both types of protection are designed to interoperate with other vendor equipment.

Such functionality makes link protection and node-link protection excellent choices for scalability, redundancy, and performance in MPLS-enabled networks.

Related Information For additional information about MPLS fast reroute and MPLS protection methods, see the following:

- *Junos Feature Guide*
- *Junos MPLS Applications Configuration Guide*
- Semeria, Chuck. *RSVP Signaling Extensions for MPLS Traffic Engineering*. White paper. 2002
- Semeria, Chuck. *IP Dependability: Network Link and Node Protection*. White paper. 2002
- RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

Verify That Link Protection Is Up

Purpose When you verify link protection, you must check that the bypass LSP is up. You can also check the number of LSPs protected by the bypass. In the network shown in *Many-to-One or Link Protection*, a bypass path should be up to protect the link between **R1** and **R2**, or next-hop **10.0.12.14**, and the two LSPs traversing the link, **lsp2-r1-to-r5** and **lsp1-r6-to-r0**.

Action To verify link protection (many-to-one backup), enter the following Junos OS CLI operational mode commands on the ingress router:

```
user@host> show mpls lsp extensive
user@host> show rsvp session detail
user@host> show rsvp interface
```

Sample Output

```
user@R1> show mpls lsp extensive | no-more
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: lsp2-r1-to-r5
  ActivePath: via-r2 (primary)
  Link protection desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary via-r2 State: Up
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
      10.0.12.14(Label=101264) 10.0.24.2(Label=100736) 10.0.45.2(Label=3)
    6 Jun 16 14:06:33 Link-protection Up
    5 Jun 16 14:05:39 Selected as active path
    4 Jun 16 14:05:39 Record Route: 10.0.12.14(Label=101264)
10.0.24.2(Label=100736) 10.0.45.2(Label=3)
    3 Jun 16 14:05:39 Up
    2 Jun 16 14:05:39 Originate Call
    1 Jun 16 14:05:39 CSPF: computation result accepted
  Created: Fri Jun 16 14:05:38 2006
Total 1 displayed, Up 1, Down 0
```

[...Output truncated...]

Transit LSP: 2 sessions

192.168.0.1

From: 192.168.6.1, LSPstate: Up, ActiveRoute: 0

LSPname: lsp1-r6-to-r0, LSPpath: Primary

Suggested label received: -, Suggested label sent: -

Recovery label received: -, Recovery label sent: 101296

Resv style: 1 SE, Label in: 100192, Label out: 101296

Time left: 116, Since: Mon Jun 19 10:26:32 2006

Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500

Port number: sender 1 receiver 58739 protocol 0

Link protection desired

Type: Link protected LSP, using Bypass->10.0.12.14

1 Jun 19 10:26:32 Link protection up, using Bypass->10.0.12.14

PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 579 pkts

Adspec: received MTU 1500 sent MTU 1500

PATH sentto: 10.0.12.14 (fe-0/1/0.0) 474 pkts

RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 501 pkts

Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

[...Output truncated...]

Meaning The sample output from ingress router R1 shows that **lsp2-r1-to-r5** and **lsp1-r6-to-r0** have link protection up, and both LSPs are using the bypass path, **10.0.12.14**. However, the **show mpls lsp** command does not list the bypass path. For information about the bypass path, use the **show rsvp session** command.

```

Sample Output user@R1> show rsvp session detail
Ingress RSVP: 2 sessions
192.168.2.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->10.0.12.14
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101456
  Resv style: 1 SE, Label in: -, Label out: 101456
  Time left: -, Since: Fri May 26 18:38:09 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 18709 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 2
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.17.14 (fe-0/1/1.0) 51939 pkts
  RESV rcvfrom: 10.0.17.14 (fe-0/1/1.0) 55095 pkts
  Explct route: 10.0.17.14 10.0.27.1
  Record route: <self> 10.0.17.14 10.0.27.1

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp2-r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101264
  Resv style: 1 SE, Label in: -, Label out: 101264
  Time left: -, Since: Fri Jun 16 14:05:39 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 18724 protocol 0
  Link protection desired
  Type: Link protected LSP
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 8477 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 8992 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.45.2
  Record route: <self> 10.0.12.14 10.0.24.2 10.0.45.2
Total 2 displayed, Up 2, Down 0

Egress RSVP: 1 sessions
192.168.1.1
  From: 192.168.5.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r5-to-r1, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 159, Since: Mon May 22 22:08:16 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 64449 protocol 0
  PATH rcvfrom: 10.0.17.14 (fe-0/1/1.0) 63145 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.59.1 10.0.79.2 10.0.17.14 <self>
Total 1 displayed, Up 1, Down 0
Transit RSVP: 2 sessions

```



```

192.168.0.1
  From: 192.168.6.1, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp1-r6-to-r0, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 101296
  Resv style: 1 SE, Label in: 100192, Label out: 101296
  Time left: 129, Since: Mon Jun 19 10:26:32 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 58739 protocol 0
  Link protection desired
  Type: Link protected LSP
  PATH rcvfrom: 10.0.16.2 (so-0/0/3.0) 3128 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 2533 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 2685 pkts
  Explct route: 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2
  Record route: 10.0.16.2 <self> 10.0.12.14 10.0.24.2 10.0.45.2 10.0.50.2

192.168.6.1
  From: 192.168.0.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r0-to-r6, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100128, Label out: 3
  Time left: 143, Since: Thu May 25 12:30:26 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 4111 protocol 0
  PATH rcvfrom: 10.0.17.14 (fe-0/1/1.0) 57716 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.16.2 (so-0/0/3.0) 54524 pkts
  RESV rcvfrom: 10.0.16.2 (so-0/0/3.0) 50534 pkts
  Explct route: 10.0.16.2
  Record route: 10.0.50.2 10.0.59.1 10.0.79.2 10.0.17.14 <self> 10.0.16.2
Total 2 displayed, Up 2, Down 0

```

Meaning The sample output from ingress router **R1** shows the ingress, egress, and transit LSPs for **R1**. Some information is similar to that found in the **show mpls lsp** command. However, because link protection is an RSVP feature, information about bypass paths is provided. The bypass path appears as a separate RSVP ingress session for the protected interface, as indicated by the **Type** field.

The bypass path name is automatically generated. By default, the name appears as **Bypass > interface-address**, where the interface address is the next downstream router's interface (**10.0.12.14**). The explicit route **10.0.17.14 10.0.27.1** for the session shows **R7** as the transit node and **R2** as the egress node.

Within the ingress RSVP section of the output, the LSP originating at **R1** (**lsp2-r1-to-r5**) is shown requesting link protection. Since a bypass path is in place to protect the downstream link, **lsp2-r1-to-r5** is associated with the bypass, as indicated by the **Link protected LSP** field.

The egress section of the output shows the return LSP **r5-to-r1**, which is not protected.

The transit section of the output shows link protection requested by **lsp1-r6-to-r0**. Since a bypass path is in place to protect the downstream link, **lsp1-r6-to-r0** is associated with the bypass, as indicated by the **Link protected LSP** field. Also included in the transit section of the output is the return LSP **r0-to-r6**, which is not protected.

Sample Output `user@R1> show rsvp interface`

```

RSVP interface: 4 active
      Active Subscr- Static      Available      Reserved      Highwater
Interface State resv  iption  BW      BW      BW      mark
fe-0/1/0.0 Up      2    100% 100Mbps 100Mbps 0bps    35Mbps
fe-0/1/1.0 Up      1    100% 100Mbps 100Mbps 0bps    0bps
fe-0/1/2.0 Up      0    100% 100Mbps 100Mbps 0bps    0bps
so-0/0/3.0 Up      1    100% 155.52Mbps 155.52Mbps 0bps    0bps

```

Meaning The sample output from ingress router **R1** shows the number of LSPs going through the interfaces configured on **R1**. The **Active resv** field shows the number of LSPs for each interface. For example, interface **fe-0/1/0.0** between **R1** and **R2** has two active reservations, **lsp1-r6-to-r0** and **lsp2-r1-to-r5**; interface **fe-0/1/1.0** between **R1** and **R7** has one, the bypass (**10.0.12.14**); interface **fe-0/1/2.0** between **R6** and **R1** has no LSP reservations; and interface **so-0/0/3.0** between **R6** and **R1** has one LSP reservation, **lsp1-r6-to-r0**.

Verify One-to-One Backup

Purpose You can verify that one-to-one backup is established by examining the ingress router and the other routers in the network.

Action To verify one-to-one backup, enter the following Junos OS CLI operational mode commands:

```

user@host> show mpls lsp ingress extensive
user@host> show rsvp session

```

Sample Output

The following sample output is from the ingress router **R1** in the network shown in *One-to-One Backup Detours*:

```

user@R1> show mpls lsp ingress extensive
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r5
  ActivePath: via-r2 (primary)
  FastReroute desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary via-r2 State: Up
  SmartOptimizeTimer: 180
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
  10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
  10=SoftPreempt):
    10.0.12.14(flag=9) 10.0.24.2(flag=1) 10.0.45.2
    8 May 11 14:51:46 Fast-reroute Detour Up
    7 May 11 14:50:55 Record Route: 10.0.12.14(flag=9) 10.0.24.2(flag=1) 10.0.45.2
    6 May 11 14:50:55 Record Route: 10.0.12.14(flag=9) 10.0.24.2 10.0.45.2
    5 May 11 14:50:52 Selected as active path
    4 May 11 14:50:52 Record Route: 10.0.12.14 10.0.24.2 10.0.45.2
    3 May 11 14:50:52 Up
    2 May 11 14:50:52 Originate Call

```

```
1 May 11 14:50:52 CSPF: computation result accepted
Created: Thu May 11 14:50:52 2006
Total 1 displayed, Up 1, Down 0
```

Meaning The sample output from **R1** shows that the **FastReroute desired** object was included in the Path messages for the LSP, allowing **R1** to select the active path for the LSP and establish a detour path to avoid **R2**.

In line 8, **Fast-reroute Detour Up** shows that the detour is operational. Lines 6 and 7 indicate that transit routers **R2** and **R4** have established their detour paths.

R2, 10.0.12.14, includes (**flag=9**), indicating that node protection is available for the downstream node and link. **R4, 10.0.24.2**, includes (**flag=1**), indicating that link protection is available for the next downstream link. In this case, **R4** can protect only the downstream link because the node is the egress router **R5**, which cannot be protected. For more information about flags, see the *Junos Feature Guide*.

The output for the **show mpls lsp extensive** command does not show the actual path of the detour. To see the actual links used by the detour paths, you must use the **show RSVP session ingress detail** command.

Sample Output The following sample output is from the ingress router **R1** in the network shown in *One-to-One Backup Detours*.

```
user@R1> show rsvp session ingress detail
Ingress RSVP: 1 sessions

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100848
  Resv style: 1 FF, Label in: -, Label out: 100848
  Time left: -, Since: Thu May 11 14:17:15 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 9228 protocol 0
  FastReroute desired
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 10.0.12.14 (fe-0/1/0.0) 35 pkts
  RESV rcvfrom: 10.0.12.14 (fe-0/1/0.0) 25 pkts
  Explt route: 10.0.12.14 10.0.24.2 10.0.45.2
  Record route: <self> 10.0.12.14 10.0.24.2 10.0.45.2
  Detour is Up
  Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Detour adspec: sent MTU 1500
  Path MTU: received 1500
  Detour PATH sentto: 10.0.17.14 (fe-0/1/1.0) 23 pkts
  Detour RESV rcvfrom: 10.0.17.14 (fe-0/1/1.0) 20 pkts
  Detour Explt route: 10.0.17.14 10.0.79.2 10.0.59.1
  Detour Record route: <self> 10.0.17.14 10.0.79.2 10.0.59.1
  Detour Label out: 100848
Total 1 displayed, Up 1, Down 0
```

Meaning The sample output from **R1** shows the RSVP session of the main LSP. The detour path is established, **Detour is Up**. The physical path of the detour is displayed in **Detour Explt route**. The detour path uses **R7** and **R9** as transit routers to reach **R5**, the egress router.

Sample Output The following sample output is from the first transit router R2 in the network shown in *One-to-One Backup Detours*:

```

user@R2> show rsvp session transit detail
Transit RSVP: 1 sessions

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 1
  LSPname: r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100448
  Resv style: 1 FF, Label in: 100720, Label out: 100448
  Time left: 126, Since: Wed May 10 16:12:21 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 9216 protocol 0
  FastReroute desired
  PATH rcvfrom: 10.0.12.13 (fe-0/1/0.0) 173 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.24.2 (so-0/0/1.0) 171 pkts
  RESV rcvfrom: 10.0.24.2 (so-0/0/1.0) 169 pkts
  Explct route: 10.0.24.2 10.0.45.2
  Record route: 10.0.12.13 <self> 10.0.24.2 10.0.45.2
  Detour is Up
  Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Detour adspec: received MTU 1500 sent MTU 1500
  Path MTU: received 1500
  Detour PATH sentto: 10.0.27.2 (so-0/0/3.0) 169 pkts
  Detour RESV rcvfrom: 10.0.27.2 (so-0/0/3.0) 167 pkts
  Detour Explct route: 10.0.27.2 10.0.79.2 10.0.59.1
  Detour Record route: 10.0.12.13 <self> 10.0.27.2 10.0.79.2 10.0.59.1
  Detour Label out: 100736
Total 1 displayed, Up 1, Down 0

```

Meaning The sample output from R2 shows the detour is established (**Detour is Up**) and avoids R4, and the link connecting R4 and R5 (10.0.45.2). The detour path is through R7 (10.0.27.2) and R9 (10.0.79.2) to R5 (10.0.59.1), which is different from the explicit route for the detour from R1. R1 has the detour passing through the 10.0.17.14 link on R7, while R1 is using the 10.0.27.2 link. Both detours merge at R9 through the 10.0.79.2 link to R5 (10.0.59.1).

Sample Output The following sample output is from the second transit router R4 in the network shown in *One-to-One Backup Detours*:

```

user@R4> show rsvp session transit detail
Transit RSVP: 1 sessions

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 1
    LSPname: r1-to-r5, LSPpath: Primary
    Suggested label received: -, Suggested label sent: -
    Recovery label received: -, Recovery label sent: 3
    Resv style: 1 FF, Label in: 100448, Label out: 3
    Time left: 155, Since: Wed May 10 16:15:38 2006
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Port number: sender 5 receiver 9216 protocol 0
    FastReroute desired
    PATH rcvfrom: 10.0.24.1 (so-0/0/1.0) 178 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 10.0.45.2 (so-0/0/2.0) 178 pkts
    RESV rcvfrom: 10.0.45.2 (so-0/0/2.0) 175 pkts
    Explct route: 10.0.45.2
    Record route: 10.0.12.13 10.0.24.1 <self> 10.0.45.2
    Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Detour adspec: received MTU 1500 sent MTU 1500
    Path MTU: received 1500
    Detour PATH sentto: 10.0.49.2 (so-0/0/3.0) 176 pkts
    Detour RESV rcvfrom: 10.0.49.2 (so-0/0/3.0) 175 pkts
    Detour Explct route: 10.0.49.2 10.0.59.1
    Detour Record route: 10.0.12.13 10.0.24.1 <self> 10.0.49.2 10.0.59.1
    Detour Label out: 100352
  Total 1 displayed, Up 1, Down 0

```

Meaning The sample output from R4 shows the detour is established (**Detour is Up**) and avoids the link connecting R4 and R5 (10.0.45.2). The detour path is through R9 (10.0.49.2) to R5 (10.0.59.1). Some of the information is similar to that found in the output for R1 and R2. However, the explicit route for the detour is different, going through the link connecting R4 and R9 (so-0/0/3 or 10.0.49.2).

Sample Output The following sample output is from **R7**, which is used in the detour path in the network shown in *One-to-One Backup Detours*:

```
user@R7> show rsvp session transit detail
Transit RSVP: 1 sessions, 1 detours

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 1
  LSPname: r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100368
  Resv style: 1 FF, Label in: 100736, Label out: 100368
  Time left: 135, Since: Wed May 10 16:14:42 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 9216 protocol 0
  Detour branch from 10.0.27.1, to skip 192.168.4.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 10.0.27.1 (so-0/0/3.0) 179 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 10.0.79.2 (so-0/0/1.0) 177 pkts
    RESV rcvfrom: 10.0.79.2 (so-0/0/1.0) 179 pkts
    Explct route: 10.0.79.2 10.0.59.1
    Record route: 10.0.12.13 10.0.27.1 <self> 10.0.79.2 10.0.59.1
    Label in: 100736, Label out: 100368
  Detour branch from 10.0.17.13, to skip 192.168.2.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 10.0.17.13 (fe-0/1/1.0) 179 pkts
    Adspec: received MTU 1500
    PATH sentto: 10.0.79.2 (so-0/0/1.0) 0 pkts
    RESV rcvfrom: 10.0.79.2 (so-0/0/1.0) 0 pkts
    Explct route: 10.0.79.2 10.0.59.1
    Record route: 10.0.17.13 <self> 10.0.79.2 10.0.59.1
    Label in: 100752, Label out: 100368
Total 1 displayed, Up 1, Down 0
```

Meaning The sample output from **R7** shows the same information as for a regular transit router used in the primary path of the LSP: the ingress address (**192.168.1.1**), the egress address (**192.168.5.1**), and the name of the LSP (**r1-to-r5**). Two detour paths are displayed; the first to avoid **R4** (**192.168.4.1**) and the second to avoid **R2** (**192.168.2.1**). Because **R7** is used as a transit router by **R2** and **R4**, **R7** can merge the detour paths together as indicated by the identical **Label out** value (**100368**) for both detour paths. Whether **R7** receives traffic from **R4** with a label value of **100736** or from **R2** with a label value of **100752**, **R7** forwards the packet to **R5** with a label value of **100368**.

Sample Output The following sample output is from R9, which is a router used in the detour path in the network shown in *One-to-One Backup Detours*:

```

user@R9> show rsvp session transit detail
Transit RSVP: 1 sessions, 1 detours

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 1
  LSPname: r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: 100352, Label out: 3
  Time left: 141, Since: Wed May 10 16:16:40 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 5 receiver 9216 protocol 0
  Detour branch from 10.0.49.1, to skip 192.168.5.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 10.0.49.1 (so-0/0/3.0) 183 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 10.0.59.1 (so-0/0/0.0) 182 pkts
    RESV rcvfrom: 10.0.59.1 (so-0/0/0.0) 183 pkts
    Explct route: 10.0.59.1
    Record route: 10.0.12.13 10.0.24.1 10.0.49.1 <self> 10.0.59.1
    Label in: 100352, Label out: 3
  Detour branch from 10.0.27.1, to skip 192.168.4.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
  Detour branch from 10.0.17.13, to skip 192.168.2.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 10.0.79.1 (so-0/0/1.0) 181 pkts
    Adspec: received MTU 1500
    PATH sentto: 10.0.59.1 (so-0/0/0.0) 0 pkts
    RESV rcvfrom: 10.0.59.1 (so-0/0/0.0) 0 pkts
    Explct route: 10.0.59.1
    Record route: 10.0.12.13 10.0.27.1 10.0.79.1 <self> 10.0.59.1
    Label in: 100368, Label out: 3
Total 1 displayed, Up 1, Down 0

```

Meaning The sample output from R9 shows that R9 is the penultimate router for the detour path, the explicit route includes only the egress link address (10.0.59.1), and the Label out value (3) indicates that R9 has performed penultimate-hop label popping. Also, the detour branch from 10.0.27.1 does not include path information because R7 has merged the detour paths from R2 and R4. Notice that the Label out value in the detour branch from 10.0.17.13 is 100368, the same value as the Label out value on R7.

Sample Output The following sample output is from the egress router R5 in the network shown in *One-to-One Backup Detours*:

```
user@R5> show rsvp session egress detail
Egress RSVP: 1 sessions, 1 detours

192.168.5.1
  From: 192.168.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r1-to-r5, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 119, Since: Thu May 11 14:44:31 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 9230 protocol 0
  FastReroute desired
  PATH rcvfrom: 10.0.45.1 (so-0/0/2.0) 258 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.12.13 10.0.24.1 10.0.45.1 <self>
  Detour branch from 10.0.49.1, to skip 192.168.5.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
  Detour branch from 10.0.27.1, to skip 192.168.4.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
  Detour branch from 10.0.17.13, to skip 192.168.2.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
  PATH rcvfrom: 10.0.59.2 (so-0/0/0.0) 254 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.12.13 10.0.24.1 10.0.49.1 10.0.59.2 <self>
  Label in: 3, Label out: -
Total 1 displayed, Up 1, Down 0
```

Meaning The sample output from R5 shows the main LSP in the **Record route** field and the detours through the network.

Verify That the Primary Path Is Operational

Purpose Primary paths must always be used in the network if they are available, therefore an LSP always moves back to the primary path after a failure, unless the configuration is adjusted. For more information on adjusting the configuration to prevent a failed primary path from reestablishing, see *Preventing Use of a Path That Previously Failed*.

Action To verify that the primary path is operational, enter the following Junos OS command-line interface (CLI) operational mode commands:

```
user@host> show mpls lsp extensive ingress
user@host> show rsvp interface
```

Sample Output 1

```

user@R1> show mpls lsp extensive ingress
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r5
  ActivePath: via-r2 (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary via-r2          State: Up
    Priorities: 6 6
    Bandwidth: 35Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 11)
  10.0.12.14 S 10.0.24.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node
  10=SoftPreempt):
      10.0.12.14 10.0.24.2
    5 Apr 29 14:40:43 Selected as active path
    4 Apr 29 14:40:43 Record Route: 10.0.12.14 10.0.24.2
    3 Apr 29 14:40:43 Up
    2 Apr 29 14:40:43 Originate Call
    1 Apr 29 14:40:43 CSPF: computation result accepted
  Standby via-r7          State: Dn
    SmartOptimizeTimer: 180
    No computed ERO.
  Created: Sat Apr 29 14:40:43 2006
Total 1 displayed, Up 1, Down 0

```

Sample Output 2

```

user@R1> show rsvp interface
RSVP interface: 3 active

```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
fe-0/1/0.0	Up	2	100%	100Mbps	100Mbps	0bps	0bps
fe-0/1/1.0	Up	1	100%	100Mbps	100Mbps	0bps	0bps
so-0/0/3.0	Up	1	100%	155.52Mbps	155.52Mbps	0bps	0bps

Meaning Sample output 1 shows that the LSP is operational and is using the primary path (**via-r2**) with **R2 (10.0.12.14)** and **R4 (10.0.24.2)** as transit routers. The priority values are the same for setup and hold, **6 6**. Priority 0 is the highest (best) priority and 7 is the lowest (worst) priority. The Junos OS default for setup and hold priority is 7:0. Unless some LSPs are more important than others, preserving the default is a good practice. Configuring a setup priority that is better than the hold priority is not allowed, resulting in a failed commit in order to avoid preemption loops.

Verify That the Secondary Path Is Established

Purpose When the secondary path is configured with the **standby** statement, the secondary path should be *up* but *not active*; it will become active if the primary path fails. A secondary path configured without the **standby** statement will not come up unless the primary path fails. To test that the secondary path is correctly configured and would come up if the

primary path were to fail, you must deactivate a link or node critical to the primary path, then issue the **show mpls lsp *lsp-path-name* extensive** command.

Action To verify that the secondary path is established, enter the following Junos OS CLI operational mode command:

Sample Output

```
user@R1> show mpls lsp extensive
```

Sample Output

The following sample output shows a correctly configured secondary path before and after it comes up. In the example, interface **fe-0/1/0** on **R2** is deactivated, which brings down the primary path **via-r2**. The ingress router **R1** switches traffic to the secondary path **via-r7**.

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.5.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r5
  ActivePath: via-r2 (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary via-r2          State: Up
    Priorities: 6 6
    Bandwidth: 35Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
10.0.12.14 S 10.0.24.2 S 10.0.45.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    10.0.12.14 10.0.24.2 10.0.45.2
    5 Apr 29 14:40:43 Selected as active path
    4 Apr 29 14:40:43 Record Route: 10.0.12.14 10.0.24.2
    3 Apr 29 14:40:43 Up
    2 Apr 29 14:40:43 Originate Call
    1 Apr 29 14:40:43 CSPF: computation result accepted
  Secondary via-r7          State: Dn
    SmartOptimizeTimer: 180
    No computed ERO.
  Created: Sat Apr 29 14:40:43 2006
Total 1 displayed, Up 1, Down 0

[edit interfaces]
user@R2# deactivate fe-0/1/0

[edit interfaces]
user@R2# show
inactive: fe-0/1/0 {
  unit 0 {
    family inet {
      address 10.0.12.14/30;
    }
    family iso;
    family mpls;
  }
}
```

```

user@R1> show mpls lsp name r1-to-r4 extensive
Ingress LSP: 1 sessions

192.168.4.1
  From: 192.168.1.1, State: Up, ActiveRoute: 0, LSPname: r1-to-r4
  ActivePath: via-r7 (secondary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary via-r2          State: Dn
    Priorities: 6 6
    Bandwidth: 35Mbps
    SmartOptimizeTimer: 180
    Will be enqueued for recomputation in 14 second(s).
  10 Apr 29 14:52:33 CSPF failed: no route toward 10.0.12.1 4[21 times]
  9 Apr 29 14:42:48 Clear Call
  8 Apr 29 14:42:48 Deselected as active
  7 Apr 29 14:42:48 Session preempted
  6 Apr 29 14:42:48 Down
  5 Apr 29 14:40:43 Selected as active path
  4 Apr 29 14:40:43 Record Route: 10.0.12.14 10.0.24.2
  3 Apr 29 14:40:43 Up
  2 Apr 29 14:40:43 Originate Call
  1 Apr 29 14:40:43 CSPF: computation result accepted
  *Standby via-r7          State: Up
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 11)
  10.0.17.14 S 10.0.47.1 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    10.0.17.14 10.0.47.1
  5 Apr 29 14:42:48 Selected as active path
  4 Apr 29 14:41:12 Record Route: 10.0.17.14 10.0.47.1
  3 Apr 29 14:41:12 Up
  2 Apr 29 14:41:12 Originate Call
  1 Apr 29 14:41:12 CSPF: computation result accepted
  Created: Sat Apr 29 14:40:43 2006
  Total 1 displayed, Up 1, Down 0

```

Meaning The sample output from egress router **R1** shows a correctly configured standby secondary path in a down state because the primary path is still up. Upon deactivation of an interface (**interface fe-0/1/0** on **R2**) critical to the primary path, the primary path **via-r2** goes down and the standby secondary path **via-r7** comes up, allowing **R1** to switch traffic to the standby secondary path.

Checklist for Checking the MPLS Layer

Problem **Description:** This checklist provides the steps and commands for checking the Multiprotocol Label Switching (MPLS) layer of the layered MPLS model. The checklist provides links to an overview of the MPLS layer and more detailed information about the commands used to investigate the problem.

[Table 18 on page 783](#) provides commands for checking the MPLS layer.

Table 18: Checklist for Checking the MPLS Layer

Tasks	Command or Action
“Checking the MPLS Layer” on page 783	
1. Verify the LSP on page 785	<pre>show mpls lsp show mpls lsp extensive show mpls lsp name <i>name</i> show mpls lsp name <i>name</i> extensive</pre>
2. Verify the LSP Route on the Transit Router on page 788	<pre>show route table mpls.0</pre>
3. Verify the LSP Route on the Ingress Router on page 789	<pre>show route <i>destination</i></pre>
4. Verify MPLS Labels with the traceroute Command on page 790	<pre>traceroute <i>hostname</i></pre>
5. Verify MPLS Labels with the ping Command on page 791	On the ingress router: <pre>ping mpls rsvp <i>lsp-name</i> detail</pre>
6. Verify the MPLS Configuration on page 792	<pre>show configuration protocols mpls show configuration interfaces</pre>
7. Take Appropriate Action on page 794	<p>The following sequence of commands addresses the specific problem described in this topic:</p> <pre>edit edit protocols mpls [edit protocols mpls] show activate interface so-0/0/3.0 show commit</pre>
8. Verify the LSP Again on page 795	<pre>show mpls lsp extensive</pre>

Checking the MPLS Layer

Purpose After you have configured the label-switched path (LSP), issued the **show mpls lsp** command, and determined that there is an error, you might find that the error is not in the physical, data link, Internet Protocol (IP), interior gateway protocol (IGP), or Resource Reservation Protocol (RSVP) layers. Continue investigating the problem at the MPLS layer of the network.

Figure 82 on page 784 illustrates the MPLS layer of the layered MPLS model.

Figure 82: Checking the MPLS Layer

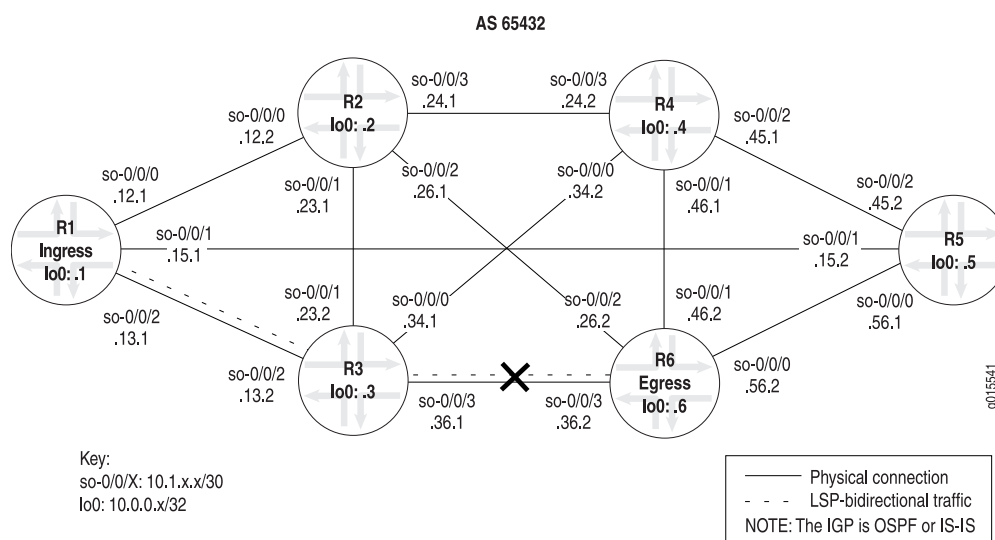
BGP Layer	traceroute <i>host-name</i> show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls.0 show route <i>address</i> traceroute <i>address</i> ping mpls rsvp <i>lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
↙ IGP and IP Layers Functioning ↘	
OSPF Layer	show ospf neighbor show configuration protocols ospf show ospf interface
IS-IS Layer	show isis adjacency show configuration protocols isis show isis interface
IP Layer	show ospf neighbor extensive show interfaces terse
IP Layer	show isis adjacency extensive show interfaces terse
Data Link Layer	show interfaces extensive <i>"JUNOS Interfaces Operations Guide"</i>
Physical Layer	show interfaces show interfaces terse ping <i>host</i>

g015547

With the MPLS layer, you check whether the LSP is up and functioning correctly. If the network is not functioning at this layer, the LSP does not work as configured.

Figure 83 on page 784 illustrates the MPLS network used in this topic.

Figure 83: MPLS Network Broken at the MPLS Layer



g015541

The network shown in Figure 83 on page 784 is a fully meshed configuration where every directly connected interface can receive and send packets to every other similar interface.

The LSP in this network is configured to run from ingress router **R1**, through transit router **R3**, to egress router **R6**. In addition, a reverse LSP is configured to run from **R6** through **R3** to **R1**, creating bidirectional traffic.

However, in this example, the reverse LSP is down without a path from **R6** to **R1**.

The cross shown in [Figure 83 on page 784](#) indicates where the LSP is broken. Some possible reasons the LSP is broken might include an incorrectly configured MPLS protocol, or interfaces that are incorrectly configured for MPLS.

In the network shown in [Figure 83 on page 784](#), a configuration error on egress router **R6** prevents the LSP from traversing the network as expected.

To check the MPLS layer, follow these steps:

1. [Verify the LSP on page 785](#)
2. [Verify the LSP Route on the Transit Router on page 788](#)
3. [Verify the LSP Route on the Ingress Router on page 789](#)
4. [Verify MPLS Labels with the traceroute Command on page 790](#)
5. [Verify MPLS Labels with the ping Command on page 791](#)
6. [Verify the MPLS Configuration on page 792](#)
7. [Take Appropriate Action on page 794](#)
8. [Verify the LSP Again on page 795](#)

Verify the LSP

Purpose Typically, you use the **show mpls lsp extensive** command to verify the LSP. However for quick verification of the LSP state, use the **show mpls lsp** command. If the LSP is down, use the **extensive** option (**show mpls lsp extensive**) as a follow-up. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action To verify that the LSP is up, enter some or all of the following commands from the ingress router:

```
user@host> show mpls lsp
user@host> show mpls lsp extensive
user@host> show mpls lsp name name
user@host> show mpls lsp name name extensive
```

Sample Output 1

```
user@R1> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
10.0.0.6    10.0.0.1    Dn     0  -              R1-to-R6
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
```

```

Total 0 displayed, Up 0, Down 0

user@R3> show mpls lsp
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R6> show mpls lsp
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
10.0.0.1    10.0.0.6    Dn     0  -              R6-to-R1
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 2

```

user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary          State: Dn
    Will be enqueued for recomputation in 22 second(s).
    1 Nov  2 14:43:38  CSPF failed: no route toward 10.0.0.6 [175 times]
    Created: Tue Nov  2 13:18:39 2004
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.1
  From: 10.0.0.6, State: Dn, ActiveRoute: 0, LSPname: R6-to-R1

```



```

ActivePath: (none)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
Primary          State: Dn
    Will be enqueued for recomputation in 13 second(s).
    1 Nov  2 14:38:12  CSPF failed: no route toward 10.0.0.1 [177 times]
Created: Tue Nov  2 13:12:22 2004
Total 1 displayed, Up 0, Down 1

```

```

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

```

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 3

```

user@R1> show mpls lsp name R1-to-R6
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPname
10.0.0.6    10.0.0.1    Dn    0  -              R1-to-R6
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 4

```

user@R1> show mpls lsp name R1-to-R6 extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Dn, ActiveRoute: 0, LSPname: R1-to-R6
  ActivePath: (none)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary          State: Dn
    Will be enqueued for recomputation in 10 second(s).
    1 Nov  2 14:51:53 CSPF failed: no route toward 10.0.0.6[192 times]
Created: Tue Nov  2 13:18:39 2004
Total 1 displayed, Up 0, Down 1

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Meaning Sample Output 1 shows a brief description of the state of the LSP for the ingress, transit, and egress routers. Output from ingress router **R1** and egress router **R6** shows that both LSPs are down, **R1-to-R6** and **R6-to-R1**. With the configured LSPs on **R1** and **R6**, we would expect egress LSP sessions on both **R1** and **R6**. In addition, transit router **R3** has no transit sessions.

Sample Output 2 shows all information about the LSPs, including all past state history and the reason why an LSP failed. Output from **R1** and **R6** indicates that there is no route to the destination because the Constrained Shortest Path First (CSPF) algorithm failed.

Sample Outputs 3 and 4 show examples of the output for the **show mpls lsp name** command with the **extensive** option. In this instance, the output is very similar to the **show mpls lsp** command because only one LSP is configured in the example network in [Figure 83 on page 784](#). However, in a large network with many LSPs configured, the results would be quite different between the two commands.

Verify the LSP Route on the Transit Router

Purpose If the LSP is up, the LSP route should appear in the **mpls.0** routing table. MPLS maintains an MPLS path routing table (**mpls.0**), which contains a list of the next label-switched router in each LSP. This routing table is used on transit routers to route packets to the next router along an LSP. If routes are not present in the output for the transit router, check the MPLS protocol configuration on the ingress and egress routers.

Action To verify the LSP route on the transit router, enter the following command from the transit router:

```
user@host> show route table mpls.0
```

Sample Output 1

```
user@R3> show route table mpls.0
mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 16w2d 21:52:40, metric 1
            Receive
1          *[MPLS/0] 16w2d 21:52:40, metric 1
            Receive
2          *[MPLS/0] 16w2d 21:52:40, metric 1
            Receive
```

Sample Output 2

```
user@R3> show route table mpls.0
mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 16w2d 22:26:08, metric 1
            Receive
1          *[MPLS/0] 16w2d 22:26:08, metric 1
            Receive
2          *[MPLS/0] 16w2d 22:26:08, metric 1
            Receive
100864     *[RSVP/7] 00:07:23, metric 1
            > via so-0/0/2.0, label-switched-path R6-to-R1
100864(S=0) *[RSVP/7] 00:07:23, metric 1
            > via so-0/0/2.0, label-switched-path R6-to-R1
100880     *[RSVP/7] 00:07:01, metric 1
            > via so-0/0/3.0, label-switched-path R1-to-R6
100880(S=0) *[RSVP/7] 00:07:01, metric 1
            > via so-0/0/3.0, label-switched-path R1-to-R6
```

Meaning Sample Output 1 from transit router **R3** shows three route entries in the form of MPLS label entries. These MPLS labels are reserved MPLS labels defined in RFC 3032, and are always present in the **mpls.0** routing table, regardless of the state of the LSP. The incoming labels assigned by RSVP to the upstream neighbor are missing from the output, indicating that the LSP is down. For more information on MPLS label entries, see *Checklist for Verifying LSP Use*.

In contrast, Sample Output 2 shows the MPLS labels and routes for a correctly configured LSP. The three reserved MPLS labels are present, and the four other entries represent the incoming labels assigned by RSVP to the upstream neighbor. These four entries represent two routes. There are two entries per route because the stack values in the MPLS header may be different. For each route, the second entry **100864 (S=0)** and **100880 (S=0)** indicates that the stack depth is not 1, and additional label values are included in the packet. In contrast, the first entry, **100864** and **100880** has an inferred **S=1** value which indicates a stack depth of 1 and makes each label the last label in that particular packet. The dual entries indicate that this is the penultimate router. For more information on MPLS label stacking, see RFC 3032, *MPLS Label Stack Encoding*.

Verify the LSP Route on the Ingress Router

Purpose Check whether the LSP route is included in the active entries in the **inet.3** routing table for the specified address.

Action To verify the LSP route, enter the following command from the ingress router:

```
user@host> show route destination
```

Sample Output 1

```
user@R1> show route 10.0.0.6
inet.0 : 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.6/32      *[IS-IS/18] 6d 01:41:37, metric 20
                  to 10.1.12.2 via so-0/0/0.0
                  > to 10.1.15.2 via so-0/0/1.0
                  to 10.1.13.2 via so-0/0/2.0

user@R6> show route 10.0.0.1
inet.0 : 28 destinations, 28 routes (28 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.1/32      *[IS-IS/18] 5d 01:01:38, metric 20
                  to 10.1.56.1 via so-0/0/0.0
                  > to 10.1.26.1 via so-0/0/2.0
                  to 10.1.36.1 via so-0/0/3.0
```

Sample Output 2

```
user@R1> show route 10.0.0.6
inet.0: 28 destinations, 28 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.6/32      *[IS-IS/18] 6d 02:13:42, metric 20
                  to 10.1.12.2 via so-0/0/0.0
                  > to 10.1.15.2 via so-0/0/1.0
```

```

to 10.1.13.2 via so-0/0/2.0

inet.3 : 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.6/32      *[RSVP/7] 00:08:07, metric 20
                 > via so-0/0/2.0, label-switched-path R1-to-R6

user@R6> show route 10.0.0.1

inet.0: 29 destinations, 29 routes (28 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.1/32      *[IS-IS/18] 5d 01:34:03, metric 20
                 to 10.1.56.1 via so-0/0/0.0
                 > to 10.1.26.1 via so-0/0/2.0
                 to 10.1.36.1 via so-0/0/3.0

inet.3 : 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.1/32      *[RSVP/7] 00:10:39, metric 20
                 > via so-0/0/3.0, label-switched-path R6-to-R1

```

Meaning Sample Output 1 shows entries in the **inet.0** routing table only. The **inet.3** routing table is missing from the output because the LSP is not working. The **inet.0** routing table is used by interior gateway protocols (IGPs) and Border Gateway Protocol (BGP) to store routing information. In this case, the IGP is Intermediate System-to-Intermediate System (IS-IS). For more information on the **inet.0** routing table, see the *Junos MPLS Applications Configuration Guide*.

If the LSP was working, we would expect to see entries that include the LSP in the **inet.3** routing table. The **inet.3** routing table is used on ingress routers to route BGP packets to the destination egress router. BGP uses the **inet.3** routing table on the ingress router to help resolve next-hop addresses. BGP is configured in the example network shown in [Figure 83 on page 784](#).

Sample Output 2 shows output you should receive when the LSP is up. The output shows both the **inet.0** and **inet.3** routing tables, indicating that LSPs **R1-to-R6** and **R6-to-R1** are available.

Verify MPLS Labels with the traceroute Command

Purpose Display the route packets take to a BGP destination where the BGP next hop for that route is the LSP egress address. By default, BGP uses the **inet.0** and the **inet.3** routing tables to resolve the next-hop address. When the next-hop address of the BGP route is not the router ID of the egress router, traffic is mapped to IGP routes, not to the LSP. Use the **traceroute** command as a debugging tool to determine whether the LSP is being used to forward traffic.

Action To verify MPLS labels, enter the following command from the ingress router:

```
user@host> traceroute hostname
```

Sample Output 1

```
user@R1> traceroute 100.100.6.1
traceroute to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
 1  10.1.12.2 (10.1.12.2)  0.627 ms  0.561 ms  0.520 ms
 2  10.1.26.2 (10.1.26.2)  0.570 ms !N  0.558 ms !N  4.879 ms !N

user@R6> traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 40 byte packets
 1  10.1.26.1 (10.1.26.1)  0.630 ms  0.545 ms  0.488 ms
 2  10.1.12.1 (10.1.12.1)  0.551 ms !N  0.557 ms !N  0.526 ms !N
```

Sample Output 2

```
user@R1> traceroute 100.100.6.1
to 100.100.6.1 (100.100.6.1), 30 hops max, 40 byte packets
 1  10.1.13.2 (10.1.13.2)  0.866 ms  0.746 ms  0.724 ms
    MPLS Label=100912 CoS=0 TTL=1 S=1
 2  10.1.36.2 (10.1.36.2)  0.577 ms !N  0.597 ms !N  0.546 ms !N

user@R6> traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 40 byte packets
 1  10.1.36.1 (10.1.36.1)  0.802 ms  0.716 ms  0.688 ms
    MPLS Label=100896 CoS=0 TTL=1 S=1
 2  10.1.13.1 (10.1.13.1)  0.570 ms !N  0.568 ms !N  0.546 ms !N
```

Meaning Sample Output 1 shows that BGP traffic is not using the LSP, consequently MPLS labels do not appear in the output. Instead of using the LSP, BGP traffic is using the IGP (IS-IS, in the example network in [Figure 83 on page 784](#)) to reach the BGP next-hop LSP egress address. The Junos OS default behavior uses LSPs for BGP traffic when the BGP next hop equals the LSP egress address.

Sample Output 2 is an example of output for a correctly configured LSP. The output shows MPLS labels, indicating that BGP traffic is using the LSP to reach the BGP next hop.

Verify MPLS Labels with the ping Command

Purpose When you ping a specific LSP, you check that echo requests are sent over the LSP as MPLS packets.

Action To verify MPLS labels, enter the following command from the ingress router to ping the egress router:

```
user@host> ping mpls rsvp lsp-name detail
```

For example:

```
user@R1> ping mpls rsvp R1-to-R6 detail
```

Sample Output 1

```
user@R1> ping mpls rsvp R1-to-R6 detail
LSP R1-to-R6 - LSP has no active path, exiting.
```

```
user@R6> ping mpls rsvp R6-to-R1 detail
LSP R6-to-R1 - LSP has no active path, exiting.
```

Sample Output 2

```
user@R1> traceroute 10.0.0.6
traceroute to 10.0.0.6 (10.0.0.6), 30 hops max, 40 byte packets
 1 10.1.15.2 (10.1.15.2) 0.708 ms 0.613 ms 0.576 ms
 2 10.0.0.6 (10.0.0.6) 0.763 ms 0.708 ms 0.700 ms
```

```
user@R1> ping mpls rsvp R1-to-R6 detail
Request for seq 1, to interface 69, label 100880
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 69, label 100880
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 69, label 100880
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 69, label 100880
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 69, label 100880
Reply for seq 5, return code: Egress-ok
```

```
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

```
user@R6> ping mpls rsvp R6-to-R1 detail
Request for seq 1, to interface 70, label 100864
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 70, label 100864
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 70, label 100864
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 70, label 100864
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 70, label 100864
Reply for seq 5, return code: Egress-ok
```

```
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

Meaning Sample Output 1 shows that the LSP does not have an active path to forward echo requests, indicating that the LSP is down.

Sample Output 2 is an example of output you should receive when the LSP is up and forwarding packets.

Verify the MPLS Configuration

Purpose After you have checked the transit and ingress routers, use the **traceroute** command to verify the BGP next hop, and used the **ping** command to verify the active path, you can check for problems with the MPLS configuration at the **[edit protocols mpls]** and **[edit interfaces]** hierarchy levels.

Action To verify the MPLS configuration, enter the following commands from the ingress, transit, and egress routers:

```
user@host> show configuration protocols mpls
```

```
user@host> show configuration interfaces
```

Sample Output 1

```
user@R1> show configuration protocols mpls
label-switched-path R1-to-R6 {
    to 10.0.0.6;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;
interface fxp0.0 {
    disable;
}

user@R3> show configuration protocols mpls
interface fxp0.0 {
    disable;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
interface so-0/0/2.0;
interface so-0/0/3.0;

user@R6> show configuration protocols mpls
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
inactive: interface so-0/0/3.0; <<< Incorrectly configured
```

Sample Output 2

```
user@R6> show configuration interfaces
so-0/0/0 {
    unit 0 {
        family inet {
            address 10.1.56.2/30;
        }
        family iso;
        family mpls;
    }
}
so-0/0/1 {
    unit 0 {
        family inet {
            address 10.1.46.2/30;
        }
        family iso;
        family mpls;
    }
}
so-0/0/2 {
    unit 0 {
        family inet {
            address 10.1.26.2/30;
        }
        family iso;
        family mpls;
    }
}
```

```

    }
  }
  so-0/0/3 {
    unit 0 {
      family inet {
        address 10.1.36.2/30;
      }
      family iso;
      family mpls;
    }
  }
  fxp0 {
    unit 0 {
      family inet {
        address 192.168.70.148/21;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.0.0.6/32;
        address 127.0.0.1/32;
      }
      family iso {
        address 49.0003.1000.0000.0006.00;
      }
    }
  }
}

```

Meaning Sample Output 1 from the ingress, transit, and egress routers shows that the configuration of interfaces on egress router **R6** is incorrect. Interface **so-0/0/3.0** is included as inactive at the **[edit protocols mpls]** hierarchy level, when it should be active because it is the interface through which the LSP travels.

Sample Output 2 shows that interfaces are correctly configured for MPLS on egress router **R6**. The interfaces are also correctly configured on the ingress and transit routers (not shown).

Take Appropriate Action

Problem **Description:** Depending on the error you encountered in your investigation, you must take the appropriate action to correct the problem. In this example, an interface is incorrectly configured at the **[edit protocols mpls]** hierarchy level on egress router **R6**.

Solution To correct the error in this example, follow these steps:

1. Activate the interface in the MPLS protocol configuration on egress router **R6**:

```

user@R6> edit
user@R6# edit protocols mpls
[edit protocols mpls]
user@R6# show
user@R6# activate interface so-0/0/3.0

```

2. Verify and commit the configuration:


```
[edit protocols mpls]
user@R6# show
user@R6# commit
```

Sample Output

```
user@R6> edit
Entering configuration mode

[edit]
user@R6# edit protocols mpls

[edit protocols mpls]
user@R6# show
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
inactive: interface so-0/0/3.0; <<< Incorrectly configured interface

[edit protocols mpls]
user@R6# activate interface so-0/0/3

[edit protocols mpls]
user@R6# show
label-switched-path R6-to-R1 {
    to 10.0.0.1;
}
inactive: interface so-0/0/0.0;
inactive: interface so-0/0/1.0;
inactive: interface so-0/0/2.0;
interface so-0/0/3.0; <<< Correctly configured interface

[edit protocols mpls]
user@R6# commit
commit complete
```

Meaning The sample output shows that the incorrectly configured interface **so-0/0/3.0** on egress router **R6** is now activated at the **[edit protocols mpls]** hierarchy level. The LSP can now come up.

Verify the LSP Again

Purpose After taking the appropriate action to correct the error, the LSP needs to be checked again to confirm that the problem in the BGP layer has been resolved.

Action To verify the LSP again, enter the following command from the ingress, transit, and egress routers:

```
user@host> show mpls lsp extensive
```

Sample Output

```
user@R1> show mpls lsp extensive
Ingress LSP: 1 sessions

10.0.0.6
```

```

From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6
ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.13.2 S 10.1.36.2 S
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    10.1.13.2 10.1.36.2
6 Nov 2 15:48:52 Selected as active path
5 Nov 2 15:48:52 Record Route: 10.1.13.2 10.1.36.2
4 Nov 2 15:48:52 Up
3 Nov 2 15:48:52 Originate Call
2 Nov 2 15:48:52 CSPF: computation result accepted
1 Nov 2 15:48:22 CSPF failed: no route toward 10.0.0.6[308 times]
Created: Tue Nov 2 13:18:39 2004
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

10.0.0.1
From: 10.0.0.6, LSPstate: Up, ActiveRoute: 0
LSPname: R6-to-R1, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 159, Since: Tue Nov 2 15:48:30 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 39106 protocol 0
PATH rcvfrom: 10.1.13.2 (so-0/0/2.0) 10 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.36.2 10.1.13.2 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

user@R3> show mpls lsp extensive
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 2 sessions

10.0.0.1
From: 10.0.0.6, LSPstate: Up, ActiveRoute: 1
LSPname: R6-to-R1, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100864, Label out: 3
Time left: 123, Since: Tue Nov 2 15:35:41 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 39106 protocol 0
PATH rcvfrom: 10.1.36.2 (so-0/0/3.0) 10 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.13.1 (so-0/0/2.0) 10 pkts

```

```

RESV rcvfrom: 10.1.13.1 (so-0/0/2.0) 10 pkts
Explct route: 10.1.13.1
Record route: 10.1.36.2 <self> 10.1.13.1

```

10.0.0.6

```

From: 10.0.0.1, LSPstate: Up, ActiveRoute: 1
LSPname: R1-to-R6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 3
Resv style: 1 FF, Label in: 100880, Label out: 3
Time left: 145, Since: Tue Nov 2 15:36:03 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 48015 protocol 0
PATH rcvfrom: 10.1.13.1 (so-0/0/2.0) 10 pkts
Adspec: received MTU 1500 sent MTU 1500
PATH sentto: 10.1.36.2 (so-0/0/3.0) 10 pkts
RESV rcvfrom: 10.1.36.2 (so-0/0/3.0) 10 pkts
Explct route: 10.1.36.2
Record route: 10.1.13.1 <self> 10.1.36.2
Total 2 displayed, Up 2, Down 0

```

```

user@R6> show mpls lsp extensive
Ingress LSP: 1 sessions

```

10.0.0.1

```

From: 10.0.0.6, State: Up, ActiveRoute: 1, LSPname: R6-to-R1
ActivePath: (primary)
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up
Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
10.1.36.1 S 10.1.13.1 S
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

10.1.36.1 10.1.13.1
6 Nov 2 15:41:44 Selected as active path
5 Nov 2 15:41:44 Record Route: 10.1.36.1 10.1.13.1
4 Nov 2 15:41:44 Up
3 Nov 2 15:41:44 Originate Call
2 Nov 2 15:41:44 CSPF: computation result accepted
1 Nov 2 15:41:14 CSPF failed: no route toward 10.0.0.1[306 times]
Created: Tue Nov 2 13:12:21 2004
Total 1 displayed, Up 1, Down 0

```

```

Egress LSP: 1 sessions

```

10.0.0.6

```

From: 10.0.0.1, LSPstate: Up, ActiveRoute: 0
LSPname: R1-to-R6, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 157, Since: Tue Nov 2 15:42:06 2004
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 48015 protocol 0
PATH rcvfrom: 10.1.36.1 (so-0/0/3.0) 11 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.1.13.1 10.1.36.1 <self>
Total 1 displayed, Up 1, Down 0

```

```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning Sample Output 1 from ingress router **R1** shows that LSP **R1-to-R6** has an active route to **R6** and the state is up.

Sample Output 2 from transit router **R3** shows that there are two transit LSP sessions, one from **R1** to **R6** and the other from **R6** to **R1**. Both LSPs are up.

Sample Output 3 from egress router **R6** shows that the LSP is up and the active route is the primary route. The LSP is now traversing the network along the expected path, from **R1** through **R3** to **R6**, and the reverse LSP, from **R6** through **R3** to **R1**.

Checklist for Working with the Layered MPLS Troubleshooting Model

Problem **Description:** This checklist provides a link to more detailed information about the layered Multiprotocol Label Switching network.

Solution [Table 19 on page 798](#) provides commands for working with the layered MPLS troubleshooting model.

Table 19: Checklist for Working with the Layered MPLS Troubleshooting Model

Tasks	Command or Action
“Understanding the Layered MPLS Troubleshooting Model” on page 798	<pre>show mpls lsp show mpls lsp extensive show mpls lsp name <i>name</i> show mpls lsp name <i>name</i> extensive</pre>

Understanding the Layered MPLS Troubleshooting Model

Problem **Description:** The layered MPLS troubleshooting model is a disciplined approach to investigating problems with an MPLS network. [Figure 84 on page 799](#) illustrates the layers in the model, and the commands you can use to structure your investigation. Because of the complexity of the MPLS network, you can obtain much better results from your investigations if you progress through the layers and verify the functioning of each layer on the ingress, egress, and transit routers before moving on to the next layer.

Solution [Figure 84 on page 799](#) shows the layered MPLS troubleshooting model that you can use to troubleshoot problems with your MPLS network.

Figure 84: Layered MPLS Network Troubleshooting Model

BGP Layer	tracroute <i>host-name</i> show bgp summary show configuration protocols bgp show route <i>destination-prefix</i> detail show route receive protocol bgp <i>neighbor-address</i>
MPLS Layer	show mpls lsp show mpls lsp extensive show route table mpls.0 show route <i>address</i> tracroute <i>address</i> ping mpls rsvp <i>lsp-name</i> detail
RSVP Layer	show rsvp session show rsvp neighbor show rsvp interface
↙ IGP and IP Layers Functioning ↘	
OSPF Layer show ospf neighbor show configuration protocols ospf show ospf interface	IS-IS Layer show isis adjacency show configuration protocols isis show isis interface
IP Layer show ospf neighbor extensive show interfaces terse	IP Layer show isis adjacency extensive show interfaces terse
Data Link Layer	show interfaces extensive <i>JUNOS Interfaces Network Operations Guide</i>
Physical Layer	show interfaces show interfaces terse ping <i>host</i>

9015528

As you move from one layer of the model to the next, you verify the correct functioning of a different component of the MPLS network and eliminate that layer as the source of the problem.

Physical Layer When you investigate the physical layer, you check that the routers are connected, and the interfaces are up and configured correctly. To check the physical layer, enter the **show interfaces**, **show interfaces terse**, and **ping** commands. If there is a problem in the physical layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command. For more information on checking the physical layer, see *Checklist for Verifying the Physical Layer*.

Data Link Layer When you investigate the data link layer, you check the encapsulation mode, for example, Point-to-Point Protocol (PPP) or Cisco High-Level Data Link Control (HDLC); PPP options, for example, header encapsulation; frame check sequence (FCS) size; and whether keepalive frames are enabled or disabled. To check the data link layer, enter the **show interfaces extensive** command. If there is a problem in the data link layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command. For more information on checking the data link layer, see *Checking the Data Link Layer* and the *Junos Interfaces Operations Guide*.

IP Layer When you investigate the IP layer, you verify that interfaces have correct IP addressing, and that the interior gateway protocol (IGP) neighbor adjacencies are established. To check the IP layer, enter the **show interfaces terse**, **show ospf neighbor extensive**, and **show isis adjacency extensive** commands. If there is a problem in the IP layer, take

appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command.

IGP Layer When you investigate the IGP layer, you verify that the the Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS) protocols are configured and running correctly. For more information about configuring OSPF and IS-IS, see *Configuring MPLS on Your Network*.

- If you have the OSPF protocol configured, you must check the IP layer first, and then the OSPF configuration. When you investigate the OSPF layer, you check that the protocol, interfaces, and traffic engineering are configured correctly. To check the OSPF layer, enter the **show configuration protocols ospf** and **show ospf interface** commands. If the problem exists in the OSPF layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command. For more information about checking the OSPF layer, see *Verifying the OSPF Protocol*.
- If you have the IS-IS protocol configured, because IS-IS and IP are independent of each other, it doesn't matter which one you check first. When you check the IS-IS configuration, you verify that IS-IS adjacencies are up, and the interfaces and IS-IS protocol are configured correctly. To check the IS-IS layer, enter the **show isis adjacency**, **show configuration protocols isis**, and **show isis interfaces** commands. If the problem exists in the IS-IS layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command. For more information about checking the IS-IS layer, see *Verifying the IS-IS Protocol*.



NOTE: The IS-IS protocol has traffic engineering enabled by default.

RSVP and MPLS Layers After you have both the IP and IGP layers functioning and the problem is still not solved, you can begin to check the Resource Reservation Protocol (RSVP) and MPLS layers to determine if the problem is in one of these layers.

- When you investigate the RSVP layer, you are checking that dynamic RSVP signaling is occurring as expected, neighbors are connected, and interfaces are configured correctly for RSVP. To check the RSVP layer, enter the **show rsvp session**, **show rsvp neighbor**, and **show rsvp interface** commands. If there is a problem in the RSVP layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command.
- When you investigate the MPLS layer, you are checking whether the LSP is up and functioning correctly. To check the MPLS layer, enter the **show mpls lsp**, **show mpls lsp extensive**, **show route table mpls.0**, **show route address**, **traceroute address**, and **ping mpls rsvp lsp-name detail** commands. If there is a problem in the MPLS layer, take appropriate action to fix it; then check that the LSP is operating as expected using the **show mpls lsp extensive** command.

BGP Layer If the problem persists after you have checked the RSVP and MPLS layers, you must verify that the Border Gateway Protocol (BGP) is working correctly. There is no point in


checking the BGP layer unless the LSP is established because BGP uses the MPLS LSP to forward traffic. When you check the BGP layer, you verify that the route is present and active, and more importantly, you ensure that the next hop is the LSP. To check the BGP layer, enter the **traceroute *host-name*, show bgp summary, show configuration protocols bgp, show route destination-prefix detail, and show route receive protocol bgp neighbor-address** commands. For more information on checking the BGP layer, see *Checking the BGP Layer*.

In reality, you could start at any level of the MPLS model to investigate a problem with your MPLS network. However, a disciplined approach, as the one described here, produces more consistent and reliable results.

Figure 85 on page 801 illustrates the basic network topology used in the following topics that demonstrate how to troubleshoot an MPLS network:

- *Checklist for Verifying the Physical Layer*
- *Checklist for Checking the Data Link Layer*
- *Checklist for Verifying the IP and IGP Layers*
- *Checklist for Checking the RSVP Layer*
- *Checklist for Checking the MPLS Layer on page 782*
- *Checklist for Checking the BGP Layer*

Figure 85: MPLS Basic Network Topology Example



The MPLS network consists of the following components:

- Router-only network with SONET interfaces
- MPLS protocol enabled on all routers, with interfaces selectively deactivated to illustrate a particular problem scenario

- All interfaces configured with MPLS
- A full-mesh IBGP topology, using AS 65432
- IS-IS or OSPF as the underlying IGP, using one level (IS-IS Level 2) or one area (OSPF area 0.0.0.0)
- A **send-statics** policy on routers R1 and R6, allowing a new route to be advertised into the network
- Two LSPs between routers R1 and R6, allowing for bidirectional traffic.

After you have configured an LSP, it is considered best practice to issue the **show mpls lsp** command to verify that the LSP is up, and to investigate further if you find an error message in the output. The error message can indicate a problem at any layer of the MPLS network.

The LSPs can be ingress, transit, or egress. Use the **show mpls lsp** command for quick verification of the LSP state, with the **extensive** option (**show mpls lsp extensive**) as a follow-up if the LSP is down. If your network has numerous LSPs, you might consider specifying the name of the LSP, using the **name** option (**show mpls lsp name *name*** or **show mpls lsp name *name* extensive**).

Action To begin the investigation of an error in your MPLS network, from the ingress router, enter some or all of the following Junos OS command-line interface (CLI) operational mode commands:

```
user@host> show mpls lsp
user@host> show mpls lsp extensive
user@host> show mpls lsp name name
user@host> show mpls lsp name name extensive
```


Sample Output 1 user@R1> show mpls lsp

Ingress LSP: 1 sessions

To	From	State	Rt	ActivePath	P	LSPname
10.0.0.6	10.0.0.1	Up	1		*	R1-to-R6

Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.0.0.1	10.0.0.6	Up	0	1 FF	3		- R6-to-R1

Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Sample Output 2 user@R1> show mpls lsp extensive

Ingress LSP: 1 sessions

10.0.0.6

From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPname: R1-to-R6

ActivePath: (primary)

LoadBalance: Random

Encoding type: Packet, Switching type: Packet, GPID: IPv4

*Primary State: Up

Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)

10.1.13.2 S 10.1.36.2 S

Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

10.1.13.2 10.1.36.2

30 Dec 28 13:47:29 Selected as active path

29 Dec 28 13:47:29 Record Route: 10.1.13.2 10.1.36.2

28 Dec 28 13:47:29 Up

27 Dec 28 13:47:29 Originate Call

26 Dec 28 13:47:29 CSPF: computation result accepted

25 Dec 28 13:46:59 CSPF failed: no route toward 10.0.0.6

24 Dec 28 13:46:39 Deselected as active

23 Dec 28 13:46:39 CSPF failed: no route toward 10.0.0.6

22 Dec 28 13:46:39 Clear Call

21 Dec 28 13:46:39 ResvTear received

20 Dec 28 13:46:39 Down

19 Dec 28 13:46:39 10.1.13.2: Session preempted

18 Dec 28 13:42:07 Selected as active path

17 Dec 28 13:42:07 Record Route: 10.1.13.2 10.1.36.2

16 Dec 28 13:42:07 Up

15 Dec 28 13:42:07 Originate Call

14 Dec 28 13:42:07 CSPF: computation result accepted

13 Dec 28 13:41:37 CSPF failed: no route toward 10.0.0.6

12 Dec 28 13:41:16 Deselected as active

11 Dec 28 13:41:16 CSPF failed: no route toward 10.0.0.6

10 Dec 28 13:41:16 Clear Call

9 Dec 28 13:41:16 ResvTear received

8 Dec 28 13:41:16 Down

7 Dec 28 13:41:16 10.1.13.2: Session preempted

6 Dec 13 11:50:15 Selected as active path

5 Dec 13 11:50:15 Record Route: 10.1.13.2 10.1.36.2

4 Dec 13 11:50:15 Up

3 Dec 13 11:50:15 Originate Call

2 Dec 13 11:50:15 CSPF: computation result accepted

1 Dec 13 11:49:45 CSPF failed: no route toward 10.0.0.6[6 times]

---(more)---[abort]

Sample Output 3

```

user@R1> show mpls lsp name R1-to-R6
Ingress LSP: 1 sessions
To          From          State Rt ActivePath      P      LSPName
10.0.0.6    10.0.0.1    Up    1
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

Sample Output 4

```

user@R1> show mpls lsp name R1-to-R6 extensive
Ingress LSP: 1 sessions

10.0.0.6
  From: 10.0.0.1, State: Up, ActiveRoute: 1, LSPName: R1-to-R6
  ActivePath: (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 20)
    10.1.13.2 S 10.1.36.2 S
      Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

        10.1.13.2 10.1.36.2
30 Dec 28 13:47:29 Selected as active path
29 Dec 28 13:47:29 Record Route: 10.1.13.2 10.1.36.2
28 Dec 28 13:47:29 Up
27 Dec 28 13:47:29 Originate Call
26 Dec 28 13:47:29 CSPF: computation result accepted
25 Dec 28 13:46:59 CSPF failed: no route toward 10.0.0.6
24 Dec 28 13:46:39 Deselected as active
23 Dec 28 13:46:39 CSPF failed: no route toward 10.0.0.6
22 Dec 28 13:46:39 Clear Call
21 Dec 28 13:46:39 ResvTear received
20 Dec 28 13:46:39 Down
19 Dec 28 13:46:39 10.1.13.2: Session preempted
18 Dec 28 13:42:07 Selected as active path
17 Dec 28 13:42:07 Record Route: 10.1.13.2 10.1.36.2
16 Dec 28 13:42:07 Up
15 Dec 28 13:42:07 Originate Call
14 Dec 28 13:42:07 CSPF: computation result accepted
13 Dec 28 13:41:37 CSPF failed: no route toward 10.0.0.6
12 Dec 28 13:41:16 Deselected as active
11 Dec 28 13:41:16 CSPF failed: no route toward 10.0.0.6
10 Dec 28 13:41:16 Clear Call
9 Dec 28 13:41:16 ResvTear received
8 Dec 28 13:41:16 Down
7 Dec 28 13:41:16 10.1.13.2: Session preempted
6 Dec 13 11:50:15 Selected as active path
5 Dec 13 11:50:15 Record Route: 10.1.13.2 10.1.36.2
4 Dec 13 11:50:15 Up
3 Dec 13 11:50:15 Originate Call
2 Dec 13 11:50:15 CSPF: computation result accepted
1 Dec 13 11:49:45 CSPF failed: no route toward 10.0.0.6[6 times]
Created: Mon Dec 13 11:47:19 2004
Total 1 displayed, Up 1, Down 0

```

Egress LSP: 1 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Meaning The sample output from the ingress router **R1** shows that the label-switched path is traversing the network as intended, from **R1** through **R3** to **R6**, and another LSP in the reverse direction, from **R6** through **R3** to **R1**.

If your network has numerous LSPs, you might consider using the **show mpls lsp** command for quick verification of the LSP state, and the **show mpls lsp name name extensive** command to continue your investigation if you find that the LSP is down.

For more information about the status and statistics of the **show mpls lsp** command, see *Checklist for Determining LSP Status*. For more information on the availability and valid use of an LSP, see *Checklist for Verifying LSP Use*.

In all the following topics, the network topology is broken at different layers of the network so that you can investigate various MPLS network problems. The problems presented are not inclusive. Instead, the problems serve to illustrate one possible process of investigation into the different layers of the troubleshooting model.

- Related Documentation**
- *Verifying the Physical Layer*
 - *Checking the Data Link Layer*
 - *Verifying the IP and IGP Layers*
 - *Checking the RSVP Layer*
 - [Checking the MPLS Layer on page 783](#)
 - *Checking the BGP Layer*

Verify That Load Balancing Is Working

Purpose After configuring load balancing, check that traffic is load-balanced equally across paths. In this section, the command output reflects the load-balancing configuration of the example network shown in *Load-Balancing Network Topology*. The **clear** commands are used to reset LSP and interface counters to zero so that the values reflect the operation of the load-balancing configuration.

Action To verify load balancing across interfaces and LSPs, use the following command on the ingress router:

```
user@host# show configuration
```

To verify load balancing across interfaces and LSPs, use the following commands on a transit router:

```
user@host# show route
user@host# show route forwarding-table
user@host# show mpls lsp statistics
```

```

user@host# monitor interface traffic
user@host# clear mpls lsp statistics
user@host# clear interface statistics

```

Sample Output

The following sample output is for the configuration on ingress router **R1**:

```

user@R1> show configuration | no-more
[...Output truncated...]
routing-options {
  [...Output truncated...]
  forwarding-table {
    export lbpp;
  }
}
[...Output truncated...]
policy-options {
  policy-statement lbpp {
    then {
      load-balance per-packet;
    }
  }
}

```

Meaning The sample output for the **show configuration** command on ingress router **R1** shows that load balancing is correctly configured with the **lbpp** policy statement. Also, the **lbpp** policy is exported into the forwarding table at the **[edit routing-options]** hierarchy level.

Sample Output The following sample output is from transit router **R2**:

```

user@R2> show route 192.168.0.1 terse

inet.0: 25 destinations, 27 routes (25 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 192.168.0.1/32   0 10      3           so-0/0/1.0
                                >so-0/0/2.0

[...Output truncated...]

```

Meaning The sample output for the **show route** command issued on transit router **R2** shows the two equal-cost paths (**so-0/0/1** and **so-0/0/2**) through the network to the loopback address to **R0 (192.168.0.1)**. Even though the right angle bracket (**>**) usually indicates the active route, in this instance it does not, as shown in the following four sample outputs.

Sample Output The following sample output is from transit router R2:

```
user@R2> monitor interface traffic

R2                               Seconds: 65                Time: 11:41:14

Interface  Link  Input packets  (pps)  Output packets  (pps)
so-0/0/0   Up    0              (0)    0              (0)
so-0/0/1   Up    126            (0)    164659         (2128)
so-0/0/2   Up    85219          (1004) 164598         (2128)
so-0/0/3   Up    0              (0)    0              (0)
fe-0/1/0   Up    328954         (4265) 85475          (1094)
fe-0/1/1   Up    0              (0)    0              (0)
fe-0/1/2   Up    0              (0)    0              (0)
fe-0/1/3   Up    0              (0)    0              (0)
[...Output truncated...]
```

Meaning The sample output for the **monitor interface traffic** command issued on transit router R2 shows that output traffic is evenly distributed across the two interfaces **so-0/0/1** and **so-0/0/2**.

Sample Output The following sample output is from transit router R2:

```
user@R2> show mpls lsp statistics
Ingress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 5 sessions
To          From          State  Packets  Bytes  LSPname
192.168.0.1 192.168.1.1   Up     87997   17951388 lsp1
192.168.0.1 192.168.1.1   Up     87997   17951388 lsp2
192.168.0.1 192.168.1.1   Up     87997   17951388 lsp3
192.168.0.1 192.168.1.1   Up     87997   17951388 lsp4
192.168.6.1 192.168.0.1   Up      0      0 r0-r1
Total 5 displayed, Up 5, Down 0
```

Meaning The sample output for the **show mpls lsp statistics** command issued on transit router R2 shows that output traffic is evenly distributed across the four LSPs configured on ingress router R6.

Sample Output The following sample output is from transit router R2:

```
user@R2> show route forwarding-table destination 10.0.90.14
Routing table: inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
10.0.90.12/30    user  0                ulst 262144 6
ucst 345 5 so-0/0/1.0
ucst 339 2 so-0/0/2.0
```

Meaning The sample output for the **show route forwarding-table destination** command issued on transit router R2 shows **ulst** in the **Type** field, which indicates that load balancing is working. The two unicast (**ucst**) entries in the **Type** field are the two next hops for the LSPs.

Sample Output The following sample output is from transit router R2:

```
user@R2> show route forwarding-table | find mpls
Routing table: mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                dscd  38    1
0                user  0                recv  37    3
1                user  0                recv  37    3
2                user  0                recv  37    3
100112           user  0                Swap 100032 so-0/0/1.0
100128           user  0                Swap 100048 so-0/0/1.0
100144           user  0 10.0.12.13         Swap 100096 fe-0/1/0.0
100160           user  0                Swap 100112 so-0/0/2.0
100176           user  0                Swap 100128 so-0/0/2.0
```

Meaning The sample output for the **show route forwarding-table | find mpls** command issued on transit router R2 shows the MPLS routing table that contains the labels received and used by this router to forward packets to the next-hop router. This routing table is used mostly on transit routers to route packets to the next router along an LSP. The first three labels in the **Destination** column (Label 0, Label 1, and Label 2) are automatically entered by MPLS when the protocol is enabled. These labels are reserved MPLS labels defined in RFC 3032. Label 0 is the IPv4 explicit null label. Label 1 is the MPLS equivalent of the IP Router Alert label, and Label 2 is the IPv6 explicit null label.

The remaining five labels in the **Destination** column are nonreserved labels that the router uses to forward traffic, and the last column **Netif**, shows the interfaces used to send the labeled traffic. For nonreserved labels, the second **Type** column shows the operation performed on matching packets. In this example, all non-reserved packets are swapped for outgoing packet labels. For example, packets with the label **100112** have their label swapped for **100032** before they are pushed out of interface **so-0/0/1.0**.

Verify the Operation of Uneven Bandwidth Load Balancing

Purpose When a router is performing unequal-cost load balancing between LSPs paths, the **show route detail** command displays a balance field associated with each next hop being used.

Action To verify that an RSVP LSP is unevenly load-balanced, use the following Junos OS CLI operational mode commands:

```
user@host> show route protocol rsvp detail
user@host> show mpls lsp statistics
```

Sample Output

```
user@R1> show route protocol rsvp detail

inet.0: 25 destinations, 25 routes (25 active, 0 holddown, 0 hidden)
10.0.90.14/32 (1 entry, 1 announced)
  State: <FlashAll>
    *RSVP Preference: 7
      Next-hop reference count: 7
      Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 10%
        Label-switched-path lsp1
        Label operation: Push 100768
      Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 20%
        Label-switched-path lsp2
        Label operation: Push 100736
      Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 30%,
selected
        Label-switched-path lsp3
        Label operation: Push 100752
      Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 40%
        Label-switched-path lsp4
        Label operation: Push 100784
      State: <Active Int>
      Local AS: 65432
      Age: 8:03 Metric: 4
      Task: RSVP
      Announcement bits (2): 0-KRT 4-Resolve tree 1
      AS path: I
inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
192.168.0.1/32 (1 entry, 1 announced)
  State: <FlashAll>
    *RSVP Preference: 7
      Next-hop reference count: 7
      Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 10%
        Label-switched-path lsp1
        Label operation: Push 100768
      Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 20%
        Label-switched-path lsp2
        Label operation: Push 100736
      Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 30%
        Label-switched-path lsp3
        Label operation: Push 100752
      Next hop: 10.0.12.14 via fe-0/1/0.0 weight 0x1 balance 40%,
selected
        Label-switched-path lsp4
        Label operation: Push 100784
      State: <Active Int>
      Local AS: 65432
      Age: 8:03 Metric: 4
      Task: RSVP
      Announcement bits (1): 1-Resolve tree 1
      AS path: I

user@R1> show mpls lsp statistics
Ingress LSP: 4 sessions
```

To	From	State	Packets	Bytes	LSPname
192.168.0.1	192.168.1.1	Up	10067	845628	lsp1
192.168.0.1	192.168.1.1	Up	20026	1682184	lsp2
192.168.0.1	192.168.1.1	Up	29796	2502864	lsp3
192.168.0.1	192.168.1.1	Up	40111	3369324	lsp4

Total 4 displayed, Up 4, Down 0

Egress LSP: 1 sessions

To	From	State	Packets	Bytes	LSPname
192.168.1.1	192.168.0.1	Up	NA	NA	r0-r1

Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions

Total 0 displayed, Up 0, Down 0

Meaning The sample output from ingress router **R1** shows that traffic is distributed according to the LSP bandwidth configuration, as indicated by the **Balance: xx%** field. For example, **lsp1** has 10 Mbps of bandwidth configured, as reflected in the **Balance: 10%** field.

PART 9

Configuration Statements

- [MPLS Configuration Statements on page 813](#)
- [RSVP Configuration Statements on page 967](#)
- [LDP Configuration Statements on page 1019](#)
- [CCC and TCC Configuration Statements on page 1085](#)
- [GMPLS Configuration Statements on page 1101](#)
- [PCEP Configuration Statements on page 1125](#)

CHAPTER 23

MPLS Configuration Statements

- [\[edit protocols mpls\] Hierarchy Level on page 817](#)
- [\[edit logical-systems\] Hierarchy Level on page 823](#)
- [\[edit protocols connections\] Hierarchy Level on page 824](#)
- [\[edit protocols link-management\] Hierarchy Level on page 824](#)
- [adaptive on page 825](#)
- [adjust-interval on page 826](#)
- [adjust-threshold on page 826](#)
- [adjust-threshold-activate-bandwidth on page 827](#)
- [adjust-threshold-overflow-limit on page 827](#)
- [adjust-threshold-underflow-limit on page 828](#)
- [admin-down on page 828](#)
- [admin-group \(for Interfaces\) on page 829](#)
- [admin-group \(for LSPs\) on page 829](#)
- [admin-group-extended on page 830](#)
- [admin-groups on page 831](#)
- [admin-groups-extended on page 832](#)
- [admin-groups-extended-range on page 833](#)
- [advertise-mode \(MPLS\) on page 834](#)
- [advertisement-hold-time on page 835](#)
- [allow-fragmentation on page 835](#)
- [always-mark-connection-protection-tlv on page 836](#)
- [associate-backup-pe-groups on page 836](#)
- [associate-lsp on page 837](#)
- [auto-bandwidth \(MPLS Tunnel\) on page 838](#)
- [auto-bandwidth \(MPLS Statistics\) on page 839](#)
- [auto-policing on page 840](#)
- [backup-pe-group on page 841](#)
- [bandwidth \(Fast Reroute, Signaled, and Multiclass LSPs\) on page 842](#)

- [bandwidth \(Static LSP\) on page 843](#)
- [bandwidth-model on page 844](#)
- [bandwidth-percent on page 845](#)
- [bfd-liveness-detection \(Protocols MPLS\) on page 846](#)
- [class-of-service \(Protocols MPLS\) on page 847](#)
- [container-label-switched-path on page 848](#)
- [corouted-bidirectional on page 849](#)
- [corouted-bidirectional-passive on page 849](#)
- [credibility on page 850](#)
- [database on page 851](#)
- [delay \(querier\) on page 852](#)
- [delay \(responder\) on page 853](#)
- [description \(Protocols MPLS\) on page 854](#)
- [deselect-on-bandwidth-failure on page 855](#)
- [diffserv-te on page 856](#)
- [disable \(Protocols MPLS\) on page 857](#)
- [dynamic-tunnels on page 858](#)
- [egress-protection \(MPLS\) on page 859](#)
- [encoding-type on page 860](#)
- [entropy-label on page 860](#)
- [ethernet-vlan \(Protocols Link Management\) on page 861](#)
- [exclude \(for Administrative Groups\) on page 861](#)
- [exclude \(for Fast Reroute\) on page 862](#)
- [exclude-srlg on page 863](#)
- [expand-loose-hop on page 864](#)
- [explicit-null \(Protocols MPLS\) on page 865](#)
- [export \(MPLS Traffic engineering database\) on page 866](#)
- [failure-action \(Protocols MPLS\) on page 867](#)
- [family mpls on page 868](#)
- [fast-reroute \(Protocols MPLS\) on page 870](#)
- [fate-sharing on page 871](#)
- [from \(Protocols MPLS\) on page 872](#)
- [gpip on page 873](#)
- [gre \(Routing Options\) on page 874](#)
- [hop-limit on page 875](#)
- [import \(MPLS Traffic Engineering Database\) on page 876](#)
- [include-all \(for Administrative Groups\) on page 877](#)

- [include-all \(for Fast Reroute\) on page 877](#)
- [include-any \(for Administrative Groups\) on page 878](#)
- [include-any \(for Fast Reroute\) on page 878](#)
- [ingress \(LSP\) on page 879](#)
- [install \(Protocols MPLS\) on page 880](#)
- [ingress-policy on page 881](#)
- [interface \(Protocols MPLS\) on page 882](#)
- [inter-domain on page 883](#)
- [ipv6-tunneling on page 883](#)
- [label-switched-path \(Protocols MPLS\) on page 884](#)
- [label-switched-path-template \(Container LSP\) on page 887](#)
- [ldp-tunneling on page 887](#)
- [least-fill on page 887](#)
- [link-protection \(Dynamic LSPs\) on page 888](#)
- [link-protection \(Static LSPs\) on page 889](#)
- [load-balance-label-capability on page 889](#)
- [log-updown \(Protocols MPLS\) on page 890](#)
- [loss \(querier\) on page 891](#)
- [loss \(responder\) on page 892](#)
- [loss-delay \(querier\) on page 893](#)
- [lsp-attributes on page 894](#)
- [maximum-bandwidth \(Protocols MPLS\) on page 894](#)
- [maximum-labels on page 895](#)
- [minimum-bandwidth-adjust-interval on page 896](#)
- [minimum-bandwidth-adjust-threshold-change on page 896](#)
- [minimum-bandwidth-adjust-threshold-value on page 897](#)
- [metric \(Protocols MPLS\) on page 898](#)
- [minimum-bandwidth on page 898](#)
- [monitor-bandwidth on page 899](#)
- [most-fill on page 899](#)
- [mpls \(Protocols\) on page 899](#)
- [mpls-tp-mode on page 900](#)
- [mtu-signaling on page 900](#)
- [next-hop \(Protocols MPLS\) on page 901](#)
- [no-bfd-triggered-local-repair on page 902](#)
- [no-cspf on page 903](#)
- [no-decrement-ttl on page 904](#)

- [no-install-to-address](#) on page 905
- [no-load-balance-label-capability](#) on page 905
- [no-mcast-replication](#) on page 906
- [no-propagate-ttl](#) on page 907
- [no-transit-statistics](#) on page 907
- [no-trap](#) on page 908
- [node-protection \(Static LSP\)](#) on page 909
- [normalization](#) on page 910
- [oam \(Protocols MPLS\)](#) on page 911
- [optimize-adaptive-teardown](#) on page 912
- [optimize-aggressive](#) on page 913
- [optimize-hold-dead-delay](#) on page 914
- [optimize-switchover-delay](#) on page 915
- [optimize-timer \(Protocols MPLS\)](#) on page 916
- [p2mp \(Protocols MPLS\)](#) on page 917
- [p2mp-lsp-next-hop](#) on page 918
- [path \(Protocols MPLS\)](#) on page 919
- [path-mtu](#) on page 920
- [per-prefix-label](#) on page 921
- [performance-monitoring \(Protocols MPLS\)](#) on page 922
- [policing \(Protocols MPLS\)](#) on page 923
- [pop](#) on page 924
- [preference \(Protocols MPLS\)](#) on page 925
- [primary \(Protocols MPLS\)](#) on page 926
- [priority \(Protocols MPLS\)](#) on page 927
- [protection-revert-time](#) on page 928
- [push](#) on page 929
- [random](#) on page 930
- [record](#) on page 931
- [retry-limit](#) on page 932
- [retry-timer](#) on page 932
- [revert-timer](#) on page 933
- [rpf-check-policy \(Routing Options\)](#) on page 934
- [rsvp-error-hold-time](#) on page 935
- [sampling \(Protocols MPLS\)](#) on page 936
- [secondary \(Protocols MPLS\)](#) on page 937
- [select](#) on page 938

- [signal-bandwidth](#) on page 938
- [smart-optimize-timer](#) on page 939
- [soft-preemption \(Protocols MPLS\)](#) on page 940
- [splitting-merging](#) on page 941
- [srlg](#) on page 942
- [srlg-cost](#) on page 943
- [srlg-value](#) on page 943
- [standby](#) on page 944
- [static-label-switched-path](#) on page 945
- [statistics \(Protocols MPLS\)](#) on page 947
- [swap](#) on page 948
- [switch-away-lsps](#) on page 949
- [switching-type](#) on page 950
- [sync-active-path-bandwidth](#) on page 951
- [te-class-matrix](#) on page 952
- [to](#) on page 953
- [traceoptions \(Protocols MPLS\)](#) on page 954
- [traffic-class \(delay\)](#) on page 956
- [traffic-class \(loss\)](#) on page 958
- [traffic-class \(loss-delay\)](#) on page 960
- [traffic-engineering \(Protocols MPLS\)](#) on page 962
- [traffic-engineering \(Protocols BGP\)](#) on page 963
- [transit-lsp-association](#) on page 964
- [ultimate-hop-popping](#) on page 965

[\[edit protocols mpls\]](#) Hierarchy Level

The following statements can also be configured at the [\[edit logical-systems *logical-system-name*\]](#) hierarchy level:

```
protocols {
  mpls {
    disable;
    admin-group {
      exclude [ group-names ];
      include-all [ group-names ];
      include-any [ group-names ];
    }
    admin-groups {
      group-name group-value;
    }
    advertisement-hold-time seconds;
    auto-policing {
      class all (drop | loss-priority-high | loss-priority-low);
    }
  }
}
```

```

    class ctnumber (drop | loss-priority-high | loss-priority-low);
  }
  bandwidth bps {
    ct0 bps;
    ct1 bps;
    ct2 bps;
    ct3 bps;
  }
  class-of-service cos-value;
  deselect-on-bandwidth-failure {
    tear-lsp;
  }
  diffserv-te {
    bandwidth-model {
      extended-mam;
      mam;
      rdm;
    }
    te-class-matrix {
      tnumber {
        priority priority;
        traffic-class ctnumber priority priority;
      }
    }
  }
  explicit-null;
  hop-limit number;
  interface (interface-name | all) {
    disable;
    admin-group [group-names];
    srlg srlg-name;
  }
  ipv6-tunneling;
  label-switched-path lsp-name {
    disable;
    adaptive;
    admin-down;
    admin-group {
      exclude [group-names];
      include-all;
      include-any [group-names];
    }
    associate-lsp;
    auto-bandwidth {
      adjust-interval seconds;
      adjust-threshold percent;
      maximum-bandwidth bps;
      minimum-bandwidth bps;
      monitor-bandwidth;
    }
    bandwidth bps {
      ct0 bps;
      ct1 bps;
      ct2 bps;
      ct3 bps;
    }
  }

```



```

class-of-service cos-value;
corouted-bidirectional;
corouted-bidirectional-passive;
description text;
entropy-label {
    ingress-policy ingress-policy-name;
}
exclude-srlg;
fast-reroute {
    (bandwidth bps | bandwidth-percent percent);
    (exclude [ group-names ] | no-exclude);
    hop-limit number;
    (include-all [ group-names ] | no-include-all);
    (include-any [ group-names ] | no-include-any);
}
from address;
hop-limit number;
inter-domain;
install {
    destination-prefix/prefix-length <active>;
}
ldp-tunneling;
least-fill;
link-protection;
lsp-attributes {
    encoding-type (ethernet | packet | pdh | sonet-sdh);
    gpid (ethernet | hdlc | ipv4 | ppp);
    signal-bandwidth type;
    switching-type (fiber | lambda | psc-1 | tdm);
}
metric number;
most-fill;
no-cspf;
no-decrement-ttl;
node-link-protection;
oam {
    bfd-liveness-detection {
        failure-action {
            make-before-break teardown-timeout seconds;
            teardown;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        minimum-transmit-interval milliseconds;
        multiplier detection-time-multiplier;
    }
    mpls-tp-mode;
    performance-monitoring {
        performance-monitoring {
            querier {
                loss {
                    traffic-class tc-value {
                        average-sample-size sample size;
                        loss-threshold loss threshold value;
                        loss-threshold-window number of samples for loss threshold;
                        measurement-quantity bytes|packets;
                    }
                }
            }
        }
    }
}

```

```

        query-interval milliseconds;
    }
}
delay {
    traffic-class tc-value {
        average-sample-size sample size;
        padding-size size;
        query-interval milliseconds;
        rtt-delay-threshold rtt threshold value;
        twcd-delay-threshold twcd threshold value;
    }
}
loss-delay {
    traffic-class tc-value {
        average-sample-size sample size;
        loss-threshold loss threshold value;
        loss-threshold-window number of samples for loss threshold;
        measurement-quantity bytes|packets;
        padding-size size;
        query-interval milliseconds;
        rtt-delay-threshold rtt threshold value;
        twcd-delay-threshold twcd threshold value;
    }
}
}
responder {
    loss {
        min-query-interval milliseconds;
    }
    delay {
        min-query-interval milliseconds;
    }
}
}
}
optimize-hold-dead-delay;
optimize-timer seconds;
p2mp path-name;
policing {
    filter filter-name;
    no-auto-policing;
}
preference preference;
primary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
}

```

```

    }
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    priority setup-priority reservation-priority;
    (record | no-record);
    select (manual | unconditional);
  }
  priority setup-priority reservation-priority;
  (random | least-fill | most-fill);
  (record | no-record);
  retry-limit number;
  retry-timer seconds;
  revert-timer seconds;
  secondary path-name {
    adaptive;
    admin-group {
      exclude [ group-names ];
      include-all [ group-names ];
      include-any [ group-names ];
    }
    bandwidth bps {
      ct0 bps;
      ct1 bps;
      ct2 bps;
      ct3 bps;
    }
  }
  class-of-service cos-value;
  hop-limit number;
  no-cspf;
  no-decrement-ttl;
  optimize-timer seconds;
  preference preference;
  priority setup-priority reservation-priority;
  (record | no-record);
  select (manual | unconditional);
  standby;
}
soft-preemption;
standby;
to address;
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
ultimate-hop-popping;
}
log-updown {
  no-trap {
    mpls-lsp-traps;
    rfc3812-traps;
  }
  (syslog | no-syslog);

```

```

trap;
trap-path-down;
trap-path-up;
}
no-cspf;
no-decrement-ttl;
no-propagate-ttl;
oam {
    lsp-ping-interval;
    lsp-ping-multiplier;
    mpls-tp-mode;
    traceoptions;
}
optimize-aggressive;
optimize-hold-dead-delay;
optimize-timer seconds;
path path-name {
    (address | hostname) <strict | loose>;
}
path-mtu {
    allow-fragmentation;
    rsvp {
        mtu-signaling;
    }
}
preference preference;
priority setup-priority reservation-priority;
(record | no-record);
revert-timer seconds;
rsvp-error-hold-time seconds;
smart-optimize-timer seconds;
standby;
static-label-switched-path lsp-name {
    bypass bypass-name {
        bandwidth bps;
        description string;
        next-hop (address | interface-name | address/interface-name);
        push out-label;
        to address;
    }
    entropy-label {
        ingress-policy ingress-policy-name;
    }
    ingress {
        bandwidth bps;
        class-of-service cos-value;
        description string;
        install {
            destination-prefix <active>;
        }
        link-protection bypass-name name;
        metric metric;
        next-hop (address | interface-name | address/interface-name);
        node-protection bypass-name name next-next-label label;
        no-install-to-address;
        policing {

```

```

        filter filter-name;
        no-auto-policing;
    }
    preference preference;
    push out-label;
    to address;
}
transit incoming-label {
    bandwidth bps;
    description string;
    link-protection bypass-name name;
    next-hop (address | interface-name | address/interface-name);
    node-protection bypass-name name next-next-label label;
    pop;
    swap out-label;
}
statistics {
    auto-bandwidth;
    file filename <files number> <size size> <world-readable | no-world-readable>;
    interval seconds;
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
}
traffic-engineering (bgp | bgp-igp | bgp-igp-both-ribs | mpls-forwarding);
transit-lsp-association;
ultimate-hop-popping;
}
}

```

[edit logical-systems] Hierarchy Level

The following MPLS protocol statements can be configured at the [edit logical-systems] hierarchy level. This is not a comprehensive list of statements available for logical systems. Only the statements that are also documented in this manual are listed here. For more information about logical systems, see the *Logical Systems Feature Guide for Routing Devices*.



NOTE: Beginning with Junos OS Release 9.3, the logical router feature has been renamed logical system.

All configuration statements, operational commands, show command outputs, error messages, log messages, and SNMP MIB objects that contain the string logical-router or logical-routers have been changed to logical-system and logical-systems, respectively.

```

logical-systems {
    logical-system-name {
        protocols {
            connections {
                connections-configuration;
            }
        }
    }
}

```

```
    }  
    ldp {  
        ldp-configuration;  
    }  
    link-management {  
        link-management-configuration;  
    }  
    mpls {  
        mpls-configuration;  
    }  
    rsvp {  
        rsvp-configuration;  
    }  
}  
}
```

[\[edit protocols connections\] Hierarchy Level](#)

The following statements can also be configured at the **[edit logical-systems *logical-system-name*]** hierarchy level:

```
protocols {  
    connections {  
        interface-switch connection-name {  
            interface interface-name.unit-number;  
        }  
        lsp-switch connection-name {  
            transmit-lsp label-switched-path;  
            receive-lsp label-switched-path;  
        }  
        p2mp-receive-switch {  
            output-interface interface-name.unit-number;  
            receive-p2mp-lsp receiving-point-to-multipoint-lsp;  
        }  
        p2mp-transmit-switch {  
            input-interface input-interface-name.unit-number;  
            transmit-p2mp-lsp transmitting-point-to-multipoint-lsp;  
        }  
        remote-interface-switch connection-name {  
            interface interface-name.unit-number;  
            transmit-lsp label-switched-path;  
            receive-lsp label-switched-path;  
        }  
    }  
}
```

[\[edit protocols link-management\] Hierarchy Level](#)

The following statements can also be configured at the **[edit logical-systems *logical-system-name*]** hierarchy level:

```
protocols {  
    link-management {  
        peer peer-name {
```

```

    address address;
    control-channel [ control-channel-interfaces ];
    te-link [ te-link-names ];
  }
  te-link te-link-name {
    disable;
    interface interface-name {
      disable;
      local-address address;
      remote-address address;
      remote-id id-number;
    }
    label-switched-path label-switched-path-name;
    local-address address;
    remote-address address;
    remote-id id-number;
  }
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}

```

adaptive

Syntax	<code>adaptive;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	During reroute, do not double-count bandwidth on links shared by the old and new paths. Including this statement causes RSVP to use shared explicit (SE) reservation styles and assists in smooth transition during rerouting.
Default	The configured object is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Adaptive LSPs on page 249

adjust-interval

Syntax	<code>adjust-interval <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Specify the bandwidth reallocation interval.
Options	<i>seconds</i> —Bandwidth reallocation interval, in seconds. Range: 300 through 315,360,000 seconds Default: 86,400 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Automatic Bandwidth Allocation Interval on page 259

adjust-threshold

Syntax	<code>adjust-threshold <i>percent</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Specify how sensitive the automatic bandwidth adjustment for a label-switched path (LSP) is to changes in bandwidth utilization.
Options	<i>percent</i> —Bandwidth demand for the current bandwidth adjustment interval is determined and compared to the LSP's current bandwidth allocation. If the percentage difference in bandwidth is greater than or equal to the percentage specified by this statement, the LSP's bandwidth is adjusted to the current bandwidth demand.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Automatic Bandwidth Adjustment Threshold on page 260

adjust-threshold-activate-bandwidth

Syntax	adjust-threshold-activate-bandwidth <i>bps</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced before Junos OS Release 14.1.
Description	Specify an absolute value to prevent automatic adjustment of signaled bandwidth and aggressive re-signaling of a label-switched path (LSP) when the actual bandwidth over the LSP is below the configured threshold, although the adjust-threshold percentage condition is satisfied.
Options	<i>bps</i> —Amount of bandwidth that is compared with the maximum of all traffic samples during an adjustment interval. If the maximum average bandwidth is less than this configured value, automatic bandwidth adjustment or re-signaling does not happen, even if the adjust-threshold percentage condition is satisfied.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Automatic Bandwidth Adjustment Threshold on page 260

adjust-threshold-overflow-limit

Syntax	adjust-threshold-overflow-limit <i>number</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Specify the number of consecutive bandwidth overflow samples before triggering a bandwidth adjustment.
Options	<i>number</i> —Number of consecutive bandwidth overflow samples. Range: 1 through 65,535 Default: This feature is disabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring a Limit on Bandwidth Overflow and Underflow Samples on page 261

adjust-threshold-underflow-limit

Syntax	adjust-threshold-underflow-limit <i>number</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced in Junos OS Release 11.3. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Specify the number of consecutive bandwidth underflow samples before triggering a bandwidth adjustment.
Options	number —Number of consecutive bandwidth underflow samples. Range: 1 through 65,535 Default: This feature is disabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Limit on Bandwidth Overflow and Underflow Samples on page 261

admin-down

Syntax	admin-down;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Set a nonpacket GMPLS LSP to the administrative down state. This statement does not affect control path setup or data forwarding for packet LSPs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Allowing Nonpacket GMPLS LSPs to Establish Paths Through Routers Running Junos OS on page 690

admin-group (for Interfaces)

Syntax	<code>admin-group [<i>group-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i>], [edit protocols mpls interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define administrative groups for an interface.
Options	<i>group-names</i> —One or more names of groups defined with the admin-groups statement at the [edit protocols mpls] hierarchy level.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Administrative Groups for LSPs on page 240 • admin-groups on page 831

admin-group (for LSPs)

Syntax	<pre>admin-group { exclude [<i>group-names</i>]; include-all [<i>group-names</i>]; include-any [<i>group-names</i>]; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define the administrative groups to include or exclude an LSP and a path's primary and secondary paths.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Administrative Groups for LSPs on page 240

admin-group-extended

Syntax	<pre>admin-group-extended { apply-groups <i>group-value</i>; apply-groups-except <i>group-value</i>; exclude [<i>group-values</i>]; include-all [<i>group-values</i>]; include-any [<i>group-values</i>]; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]</pre>
Release Information	Statement introduced in Junos OS Release 11.1.
Description	<p>Specifies the group name and group identifier for an administrative group. The group identifier must be within the range of values specified by the admin-groups-extended-range statement. The extended administrative group values are global and must be identically configured on all the supported routers participating in the network. The domain-wide extended administrative groups database, learned from other routers through IGP flooding, is used by CSPF for path computation.</p>
Options	<p>apply-groups—Apply the specified administrative groups for the LSP or for the primary and secondary paths.</p> <p>apply-groups-except—Exclude the specified administrative groups from the LSP or from the primary and secondary paths.</p> <p>exclude—Define the administrative groups to exclude from an LSP or from the primary and secondary paths.</p> <p>include-all—Require the LSP to traverse links that include all of the defined administrative groups.</p> <p>include-any—Define the administrative groups to include for an LSP for the primary and secondary paths.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Extended Administrative Groups for LSPs on page 242 • Configuring Administrative Groups for LSPs on page 240 • admin-groups-extended on page 832 • admin-groups-extended-range on page 833

admin-groups

Syntax	admin-groups { <i>group-name group-value</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure administrative groups to implement link coloring of resource classes.
Options	<p><i>group-name</i>—Name of the group. You can assign up to 32 names. The names and their corresponding values must be identical across all routers within a single domain.</p> <p><i>group-value</i>—Value assigned to the group. The names and their corresponding values must be identical across all routers within a single domain.</p> <p>Range: 0 through 31</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Administrative Groups for LSPs on page 240 • admin-group (for Interfaces) on page 829

admin-groups-extended

Syntax	<code>admin-groups-extended <i>group-name</i> { group-value <i>group-identifier</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options], [edit protocols mpls interface <i>interface-name</i>], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 11.1. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Specifies the group name and group identifier for an administrative group. The group identifier must be within the range of values specified by the admin-groups-extended-range statement. The extended administrative group values are global and must be identically configured on all the supported routers participating in the network. The domain-wide extended administrative groups database, learned from other routers through IGP flooding, is used by CSPF for path computation.
Options	<i>group-name</i> —The range of configurable values is between 32 and 4,294,967,295. The range maximum must be greater than the range minimum. <i>group-value group-identifier</i> —The group identifier must be within the range of configurable values, 32 and 4,294,967,295.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Extended Administrative Groups for LSPs on page 242• Configuring Administrative Groups for LSPs on page 240• admin-group-extended on page 830• admin-groups-extended-range on page 833

admin-groups-extended-range

Syntax	<pre>admin-groups-extended-range { maximum <i>maximum-number</i>; minimum <i>minimum-number</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 11.1. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Enables you to configure extended administrative groups, represented by a 32-bit value, expanding the number of administrative groups supported in the network beyond just 32. In MPLS traffic engineering, a link can be configured with a set of administrative groups (also known as colors or resource classes). Administrative groups are carried in IGP (OSPFv2 and IS-IS) as a 32-bit value assigned to each link. By default, Juniper Networks routers interpret this 32-bit value as a bit mask with each bit representing a group. This normally limits each network to a total of 32 distinct administrative groups (value range 0 through 31).</p> <p>The extended administrative groups configuration accepts a set of interfaces with a corresponding set of extended administrative group names. It converts the names into a set of 32-bit values and propagates this information into the IGP. The extended administrative group values are global and must be identically configured on all the supported routers participating in the network. The domain-wide extended administrative groups database, learned from other routers through IGP flooding, is used by CSPF for path computation.</p>
Options	<p>maximum <i>maximum-number</i>—The range of configurable values is between 32 and 4,294,967,295. The range maximum must be greater than the range minimum.</p> <p>minimum <i>minimum-number</i>—The range of configurable values is between 32 and 4,294,967,295. The range maximum must be greater than the range minimum.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Extended Administrative Groups for LSPs on page 242 • Configuring Administrative Groups for LSPs on page 240 • admin-group-extended on page 830

advertise-mode (MPLS)

Syntax	advertise-mode (stub-alias stub-proxy);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls egress-protection context-identifier <i>context-id</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> egress-protection context-identifier <i>context-id</i>], [edit protocols mpls egress-protection context-identifier <i>context-id</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> egress-protection context-identifier <i>context-id</i>]
Release Information	Statement introduced in Junos OS Release 13.3.
Description	<p>Configure the method for the interior gateway protocol (IGP) to advertise egress protection availability.</p> <p>Egress protection availability is advertised in the IGP. Label protocols along with CSPF use this information to do egress protection.</p>
Options	<p>stub-alias—Context identifier has an alias.</p> <p>In the alias method, the LSP end-point address has an explicit backup egress node where the backup node can be learned or configured on the penultimate hop node (PHN) of a protected LSP. With this model, the PHN of a protected LSP sets up the bypass LSP tunnel to back up the egress node by avoiding the primary egress node. This model requires a Junos OS upgrade in core nodes, but is flexible enough to support all traffic engineering constraints.</p> <p>stub-proxy—Context-identifier has a stub proxy node.</p> <p>A stub node is one that only appears at the end of an AS path, which means it does not provide transit service. In this mode, known as the virtual or proxy mode, the LSP end-point address is represented as a node with bidirectional links, with the LSP's primary egress node and backup egress node. With this representation, the penultimate hop of the LSP primary egress point can behave like a PLR in setting up a bypass tunnel to back up the egress by avoiding the primary egress node. This model has the advantage that you do not need to upgrade Junos OS on core nodes and thereby helps operators to deploy this technology.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Egress Protection for Layer 3 VPN Edge Protection Overview</i>• <i>Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP</i>

advertisement-hold-time

Syntax	<code>advertisement-hold-time <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Do not advertise when the LSP goes from up to down, for a certain period of time known as the hold time.
Options	<i>seconds</i> —Hold time, in seconds. Range: 0 through 65,535 seconds Default: 5 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Damping Advertisement of LSP State Changes on page 268

allow-fragmentation

Syntax	<code>allow-fragmentation;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls path-mtu], [edit protocols mpls path-mtu]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Allow IP packets to be fragmented before they are encapsulated in MPLS.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling Packet Fragmentation on page 485

always-mark-connection-protection-tlv

Syntax	always-mark-connection-protection-tlv;
Hierarchy Level	[edit logical-systems <i>logical-systems-name</i> protocols mpls interface <i>interface-name</i>], [edit protocols mpls interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	<p>(MX Series routers only) Enable you to switch an LSP away from a network node using a bypass LSP. This feature could be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.</p> <p>This statement marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, you then need to configure the switch-away-lsps statement.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Switching LSPs Away from a Network Node on page 479

associate-backup-pe-groups

Syntax	associate-backup-pe-groups;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	<p>Enable an LSP to monitor the status of its destination PE router. You can configure multiple backup PE router groups using the same router's address. Backup PE router groups provide ingress PE router redundancy when point-to-multipoint LSPs are configured for multicast distribution. A failure of this LSP indicates to all of the backup PE router groups that the destination PE router is down. This statement is not tied to a specific backup PE router group. It applies to all groups that are interested in the status of the LSP to the destination address.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Enabling Point-to-Point LSPs to Monitor Egress PE Routers on page 307

associate-lsp

Syntax	<code>associate-lsp <i>lsp-name</i> { from <i>from-ip-address</i>; }</code>
Hierarchy Level	<code>[edit protocols mpls label-switched-path <i>lsp-name</i> oam]</code>
Release Information	Statement introduced in Junos OS Release 12.1.
Description	Configure associated bidirectional label-switched paths (LSPs) on the two ends of an LSP for sending and receiving GAL and G-Ach OAM messages.
Options	from <i>from-ip-address</i> —(Optional) Source address for the associated LSP configuration. If omitted, this is derived from the to address of the ingress LSP configuration.
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring the MPLS Transport Profile for OAM on page 136

auto-bandwidth (MPLS Tunnel)

Syntax	<pre>auto-bandwidth { adjust-interval <i>seconds</i>; adjust-threshold <i>percent</i>; adjust-threshold-activate-bandwidth <i>bps</i> adjust-threshold-overflow-limit <i>number</i>; adjust-threshold-underflow-limit <i>number</i>; maximum-bandwidth <i>bps</i>; minimum-bandwidth <i>bps</i>; minimum-bandwidth-adjust-interval minimum-bandwidth-adjust-threshold-change minimum-bandwidth-adjust-threshold-value monitor-bandwidth; }</pre>
Hierarchy Level	[edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Allow an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Automatic Bandwidth Allocation for LSPs on page 257• request mpls lsp adjust-autobandwidth on page 1168

auto-bandwidth (MPLS Statistics)

Syntax	auto-bandwidth;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls statistics], [edit protocols mpls statistics]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Collect statistics related to automatic bandwidth.
Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Automatic Bandwidth Allocation for LSPs on page 257• Configuring MPLS to Gather Statistics on page 342• statistics on page 947


auto-policing

Syntax	<pre>auto-policing { class all (drop loss-priority-high loss-priority-low); class <i>ctnumber</i> (drop loss-priority-high loss-priority-low); }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable the automatic policing of all the MPLS LSPs on the router or logical system.
Options	<p>class all—Apply the same policer action to all the class types (ct0, ct1, ct2, and ct3).</p> <p>class <i>ctnumber</i>—Specific class type (ct0, ct1, ct2, or ct3) to which to apply a policer action.</p> <p>Policer actions—You can specify the following policer actions:</p> <p>Default: no action</p> <ul style="list-style-type: none">• drop—Drop all packets.• loss-priority-high—Set the packet loss priority (PLP) to high.• loss-priority-low—Set the PLP to low.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• policing (Protocols MPLS) on page 923• Configuring Automatic Policers on page 349v

backup-pe-group

Syntax	<code>backup-pe-group <i>group-name</i> { backups [<i>addresses</i>]; local-address <i>address</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 9.0. Statement introduced in Junos OS Release 9.5 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 12.3 for ACX Series routers. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.
Options	backups <i>addresses</i> —Specify the address of backup PE routers for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution. local-address <i>address</i> —Specify the address of the local PE router for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution. <i>pe-group-name</i> —Specify the name for the group of PE routers that provide ingress PE router redundancy for point-to-multipoint LSPs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Ingress PE Redundancy • Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs on page 307

bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)

Syntax	<pre>bandwidth <i>bps</i> { ct0 <i>bps</i>; ct1 <i>bps</i>; ct2 <i>bps</i>; ct3 <i>bps</i>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	<p>When configuring an LSP, specify the traffic rate associated with the LSP.</p> <p>When configuring fast reroute, allocate bandwidth for the reroute path. By default, no bandwidth is reserved for the rerouted path. The fast reroute bandwidth does not need to be identical to that allocated for the LSP itself.</p> <p>When configuring a multiclass LSP, use the ctnumber bandwidth statements to specify the bandwidth to be allocated for each class type.</p>
Options	<p>bps—Bandwidth, in bits per second. You can specify this as an integer value. You can also use the abbreviations k (for a thousand), m (for a million), or g (for a billion).</p> <p>Range: Any positive integer Default: 0 (no bandwidth is reserved)</p>
<div>  NOTE: On the ACX Series, <i>bps</i> is the only supported option. </div>	
	<p>ctnumber bps—Bandwidth for the specified class type, in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations k (for a thousand), m (for a million), or g (for a billion [also called a thousand million]).</p> <p>Range: Any positive integer Default: 0 (no bandwidth is reserved)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring Fast Reroute on page 226](#)
 - [Configuring the Bandwidth Value for LSPs on page 256](#)
 - [Configuring Traffic-Engineered LSPs on page 326](#)
 - [Configuring Class-Type Bandwidth Constraints for Multiclass LSPs on page 329](#)

bandwidth (Static LSP)

Syntax	<code>bandwidth <i>bps</i></code> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> bypass], [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress], [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>], [edit protocols mpls static-label-switched-path <i>lsp-name</i> bypass], [edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress], [edit protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	When configuring a static LSP, specify the traffic rate associated with the LSP.
Options	<i>bps</i> —Bandwidth, in bits per second. You can specify this as an integer value. You can also use the abbreviations k (for a thousand), m (for a million), or g (for a billion).
	Range: Any positive integer
	Default: 0 (no bandwidth is reserved)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	• Configuring Static LSPs on page 271

bandwidth-model

Syntax	<pre>bandwidth-model { extended-mam; mam; rdm; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls diffserv-te], [edit protocols mpls diffserv-te]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Configure the bandwidth model for differentiated services. Note that you cannot configure both bandwidth models at the same time.
Options	<p>extended-mam—The extended maximum allocation model (MAM) is a bandwidth model based on MAM.</p> <p>mam—The MAM is defined in RFC 4125, <i>Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering</i>.</p> <p>rdm—The Russian dolls bandwidth allocation model (RDM) is defined in RFC 4127, <i>Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering</i>. RDM makes efficient use of bandwidth by allowing the class types to share bandwidth.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Bandwidth Model on page 316

bandwidth-percent

Syntax	<code>bandwidth-percent <i>percentage</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the percentage of bandwidth to reserve for the detour path in case the primary path for a traffic engineered LSP or a multiclass LSP fails. The percentage configured indicates the percentage of the protected path's bandwidth that is reserved for the detour path.
Options	<i>percentage</i> —The percentage of the protected path's bandwidth that is reserved for the detour path.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Fast Reroute on page 226 • Configuring Fast Reroute for Traffic-Engineered LSPs on page 327 • Configuring Fast Reroute for Multiclass LSPs on page 330

bfd-liveness-detection (Protocols MPLS)

Syntax	<pre>bfd-liveness-detection { failure-action { make-before-break teardown-timeout <i>seconds</i>; teardown; } minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; minimum-transmit-interval <i>milliseconds</i>; multiplier <i>detection-time-multiplier</i>; }</pre>
Hierarchy Level	[edit protocols mpls label-switched-path <i>lsp-name</i> oam], [edit protocols mpls oam]
Release Information	Statement introduced in Junos OS Release 7.6. failure-action option added in Junos OS Release 8.5. Statement introduced in Junos OS Release 12.2 for EX Series switches.
Description	Enable Bidirectional Forwarding Detection (BFD) for all of the MPLS LSPs or for just a specific LSP.
Options	<p>minimum-interval—Minimum transmit and receive interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>minimum-receive-interval—Minimum receive interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>minimum-transmit-interval—Minimum transmit interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>multiplier—Detection time multiplier. Range: 1 through 255 Default: 3</p> <p>The failure-action statement is explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BFD for MPLS IPv4 LSPs on page 354• Configuring Bidirectional Forwarding Detection for MPLS (CLI Procedure)

class-of-service (Protocols MPLS)

Syntax	<code>class-of-service cos-value;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.</p>
Description	<p>Class-of-service (CoS) value given to all packets in the LSP.</p> <p>The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.</p>
Options	<p>cos-value—CoS value. A higher value typically corresponds to a higher level of service.</p> <p>Range: 0 through 7</p> <p>Default: If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Class of Service for MPLS LSPs on page 244 • Configuring the Ingress Router for Static LSPs on page 271 • Configuring the Intermediate (Transit) and Egress Routers for Static LSPs on page 274

container-label-switched-path

Syntax	<pre>container-label-switched-path <i>lsp-name</i> { disable; description <i>description</i>; label-switched-path-template; splitting-merging; suffix <i>string</i>; to <i>ip-address</i>; }</pre>
Hierarchy Level	[edit protocols mpls]
Release Information	Statement introduced in Junos OS Release 14.2.
Description	Configure a multi-label-switched path (LSP) tunnel between the ingress and the egress routers. The container LSP consists of several member LSPs to the same destination.
Options	<p>disable—Disable MPLS container-label-switched path.</p> <p>description <i>description</i>—Text describing the container LSP.</p> <p>suffix <i>string</i>—Suffix to generate names of member LSPs of the container LSP.</p> <p>to <i>ip-address</i>—IP address of the egress router.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

corouted-bidirectional

Syntax	corouted-bidirectional;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify that the label-switched path be established as a corouted bidirectional packet LSP. You cannot configure this statement at the same time as the corouted-bidirectional-passive statement.
Default	This statement is disabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Corouted Bidirectional LSPs on page 220 • corouted-bidirectional-passive on page 849

corouted-bidirectional-passive

Syntax	corouted-bidirectional-passive;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify that the label-switched path be a passive LSP associated with a bidirectional LSP when it is signaled at the ingress router. This passive LSP enables the MPLS application to utilize the reverse LSP. You cannot configure this statement at the same time as the corouted-bidirectional statement.
Default	This statement is disabled by default.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Corouted Bidirectional LSPs on page 220 • corouted-bidirectional on page 849

credibility

Syntax	<pre>credibility { direct; isis-level-1; isis-level-2; ospf; static; unknown; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls traffic-engineering database export], [edit protocols mpls traffic-engineering database export]
Release Information	Statement introduced in Junos OS Release 14.2.
Description	<p>Configure preference values for entries from BGP-TE to the traffic engineering database. A protocol with a higher credibility value is preferred over a protocol with a lower credibility value.</p> <p>The credibility order for the BGP-TE protocols is as follows:</p> <ul style="list-style-type: none">• Unknown—80• OSPF—81• ISIS Level 1—82• ISIS Level 2—83• Static—84• Direct—85
Options	<p>direct—Entries sourced from directly connected links.</p> <p>isis-level-1—Entries sourced from IS-IS Level 1.</p> <p>isis-level-2—Entries sourced from IS-IS Level 2.</p> <p>ospf—Entries sourced from OSPF.</p> <p>static—Entries sourced from static configuration.</p> <p>unknown—Entries sourced from unknown entities.</p> <p>Range: 0 through 512</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• traffic-engineering on page 962

database

Syntax	<pre>database { export; import; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls traffic-engineering], [edit protocols mpls traffic-engineering]
Release Information	Statement introduced in Junos OS Release 14.2.
Description	<p>Include link and node entries from the traffic engineering database into the lsdist.0 routing information base (RIB), so it gets picked up by the BGP export policy.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• traffic-engineering on page 962

delay (querier)

Syntax	<pre>delay { traffic-class tc-value { average-sample-size sample size; padding-size size; query-interval milliseconds; rtt-delay-threshold rtt threshold value; twcd-delay-threshold twcd threshold value; } }</pre>
Hierarchy Level	[edit protocols mpls oam performance-monitoring querier], [edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier], [edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring querier], [edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring querier]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	Configure delay measurement options. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Pro-Active Loss and Delay Measurements on page 448• On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview on page 50• performance-monitoring (Protocols MPLS) on page 922

delay (responder)

Syntax	delay { min-query-interval <i>milliseconds</i> ; }
Hierarchy Level	[edit protocols mpls oam performance-monitoring responder], [edit protocols mpls label-switched-path <i>lsp-name</i> oam performance-monitoring responder], [edit protocols mpls label-switched-path <i>lsp-name</i> primary <i>path-name</i> oam performance-monitoring responder], [edit protocols mpls label-switched-path <i>lsp-name</i> secondary <i>path-name</i> oam performance-monitoring responder]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	Configure delay measurement options.
Options	min-query-interval <i>milliseconds</i> —(Optional) Specify the minimum query interval that the responder supports. If the minimum query interval of the responder is greater than the query interval configured at querier, the effective message query rate will be the minimum query interval configured for the responder. Default: 10 seconds Range: 1000 through 4294967295 milliseconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Pro-Active Loss and Delay Measurements on page 448 • On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview on page 50 • performance-monitoring (Protocols MPLS) on page 922

description (Protocols MPLS)

Syntax	<code>description text;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls</code> <code>static-label-switched-path <i>lsp-name</i> bypass],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls</code> <code>static-label-switched-path <i>lsp-name</i> ingress],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls</code> <code>static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>],</code> <code>[edit protocols mpls label-switched-path <i>lsp-name</i>],</code> <code>[edit protocols mpls static-label-switched-path <i>lsp-name</i> bypass],</code> <code>[edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress],</code> <code>[edit protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
Description	Provides a textual description of the LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the show mpls lsp detail command and has no effect on the operation of the LSP.
Options	text —Provide a textual description of the LSP. The description text can be no more than 80 characters in length.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Text Description for LSPs on page 217

deselect-on-bandwidth-failure

Syntax	<code>deselect-on-bandwidth-failure { tear-lsp; }</code>
Hierarchy Level	[edit protocols mpls], [edit protocols mpls label-switched-path <i>path-name</i>]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	Deselect an active path if it does not meet the bandwidth criteria required for path selection.
Options	tear-lsp — Bring down an active path if none of the paths are able to reserve the required bandwidth.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

diffserv-te

Syntax	<pre>diffserv-te { bandwidth-model { extended-mam; mam; rdm; } te-class-matrix { tnumber { priority <i>priority</i>; traffic-class { ctnumber <i>priority priority</i>; } } } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify properties for differentiated services in traffic engineering.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Routers for DiffServ-Aware Traffic Engineering on page 315

disable (Protocols MPLS)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls interface interface-name], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path lsp-name], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path lsp-name auto-bandwidth], [edit protocols mpls], [edit protocols mpls interface interface-name], [edit protocols mpls label-switched-path lsp-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
Description	Disable the functionality of the configured object.
Default	The configured object is enabled (operational) unless explicitly disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Minimum MPLS Configuration on page 59

dynamic-tunnels

Syntax	<pre>dynamic-tunnels <i>tunnel-name</i> { destination-networks <i>prefix</i>; gre; rsvp-te <i>entry-name</i> { destination-networks <i>network-prefix</i>; label-switched-path-template (Multicast) { default-template; template-name; } } source-address <i>address</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure a dynamic tunnel between two PE routers.
Options	<i>tunnel-name</i> —Name of the dynamic tunnel. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks</i>• <i>Configuring GRE Tunnels for Layer 3 VPNs</i>• <i>Configuring Dynamic Tunnels</i>

egress-protection (MPLS)

Syntax	<pre>egress-protection { context-identifier <i>context-id</i> { primary protector; metric <i>igp-metric-value</i>; advertise-mode (stub-alias stub-proxy); } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in Junos OS Release 10.4. Options primary , protector , and metric introduced in Junos OS Release 11.4R3. Option advertise-mode introduced in Junos OS Release 13.3.
Description	Enables an Edge Protection Virtual Circuit (EPVC) for the MPLS protocol.
Options	<p>context-identifier <i>context-id-ip-address</i>—(Optional) The context identifier IPv4 address.</p> <p>metric <i>igp-metric-value</i>—(Optional) The IGP metric value ranging from 2 through 16777215.</p> <p>(primary protector)—On the primary PE router, configure as type primary. On the protector PE router, configure as type protector.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Example: Configuring Egress Protection for Layer 3 VPN Services Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP

encoding-type

Syntax	encoding-type (ethernet packet pdh sonet-sdh);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes], [edit protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the encoding type of payload carried by the LSP. It can be any of the following: <ul style="list-style-type: none">• ethernet—Ethernet• packet—Packet• pdh—Plesiochronous digital hierarchy (PDH)• sonet-sdh—SONET/SDH
Default	packet
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Encoding Type on page 688

entropy-label

Syntax	entropy-label { ingress-policy <i>ingress-policy-name</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp],
Release Information	Statement introduced in Junos OS Release 14.1.
Description	Assists the transit router in load-balancing MPLS traffic across ECMP paths or Link Aggregation groups by introducing the entropy label to the MPLS label stack. The entropy label allows routers to load balance MPLS traffic by using a hash-input without the need to perform deep packet inspection. Deep packet inspection requires more of the router's processing power and is not a capability shared by all routers.
Options	The other statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Entropy Label for LSPs on page 218

ethernet-vlan (Protocols Link Management)

Syntax	ethernet-label { vlan-id-range <i>vlan-id-range</i> ; }
Hierarchy Level	[edit protocols link-management te-link <i>te-link-name</i>]
Release Information	Statement introduced in Junos OS Release 14.2.
Description	Specify the TE-link to be used for Layer 2 VLAN label-switched path (LSP).
Options	<i>vlan-id-range</i> <i>vlan-id-range</i> —Pool of VLAN IDs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

exclude (for Administrative Groups)

Syntax	exclude [<i>group-names</i>];
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> admin-group], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i> admin-group], [edit protocols mpls label-switched-path <i>lsp-name</i> admin-group], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i> admin-group]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define the administrative groups to exclude for an LSP or for a path's primary and secondary paths.
Options	<i>group-names</i> —Names of one or more groups defined with the admin-groups statement.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Administrative Groups for LSPs on page 240

exclude (for Fast Reroute)

Syntax	(exclude [<i>group-names</i>] no-exclude);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.
Description	Control exclusion of administrative groups: <ul style="list-style-type: none">• exclude—Define the administrative groups to exclude for fast reroute.• no-exclude—Disable administrative group exclusion.
Options	<i>group-names</i> —Names of one or more groups defined with the admin-groups statement.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Fast Reroute on page 226• admin-groups on page 831


exclude-srlg

Syntax	exclude-srlg;
Hierarchy Level	<p>[edit protocols mpls],</p> <p>[edit logical-systems logical-system-name protocols mpls],</p> <p>[edit protocols mpls label-switched-path <i>path-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>path-name</i>],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>destination</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>destination</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	<p>Exclude Shared Risk Link Group (SRLG) links for the secondary path for critical links where it is imperative to keep the secondary and primary label-switched paths completely disjoint from any common SRLG.</p> <p>When specified, the Constrained Shortest Path First (CSPF) algorithm excludes any link belonging to the set of SRLGs in the primary path. When not specified and if a link belongs to the set of SRLGs in the primary path, CSPF adds the SRLG cost to the metric, but still accepts the link for computing the path.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Excluding SRLG Links Completely for the Secondary LSP on page 90

expand-loose-hop

Syntax	expand-loose-hop;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced in Junos OS Release 7.6. Point-to-multipoint LSP support introduced in Junos OS Release 11.2.
Description	<p>Allow an LSP to traverse multiple OSPF areas within a service provider's network.</p> <p>Allows a point-to-multipoint LSP to span multiple domains in a network. Effectively, this allows you to configure one or more sub-LSPs (branches) in separate network domains. Examples of such domains include OSPF areas and autonomous systems (ASs). A sub-LSP of an inter-domain point-to-multipoint LSP can be intra-area, inter-area, or inter-AS, depending on the location of the egress node (leaf) with respect to the ingress node (source). Only OSPF areas are supported for inter-domain point-to-multipoint LSPs. IS-IS levels are not supported.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling Interarea Traffic Engineering on page 339• Configuring Inter-Domain Point-to-Multipoint LSPs on page 303

explicit-null (Protocols MPLS)

Syntax	explicit-null;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
Description	Advertise label 0 to the egress router of an LSP.
Default	If you do not include the explicit-null statement in the MPLS configuration, label 3 (implicit null) is advertised.
<div>  <p>NOTE: Junos OS does not support explicit null routes with next hops to virtual tunnel (vt-) interfaces.</p> </div>	
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring RSVP to Pop the Label on the Ultimate-Hop Router on page 489

export (MPLS Traffic engineering database)

Syntax	<pre>export { credibility; policy <i>policy-name</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls traffic-engineering database], [edit protocols mpls traffic-engineering database]
Release Information	Statement introduced in Junos OS Release 14.2.
Description	Configure the traffic engineering database export-related parameters.
Options	<p>policy <i>policy-name</i>—Name of the export policy.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• traffic-engineering on page 962

failure-action (Protocols MPLS)

Syntax	<pre>failure-action { make-before-break teardown-timeout <i>seconds</i>; teardown; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols mpls oam bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> oam bfd-liveness-detection], [edit protocols mpls label-switched-path <i>lsp-name</i> oam bfd-liveness-detection], [edit protocols mpls oam bfd-liveness-detection]</pre>
Release Information	Statement introduced in Junos OS Release 9.4.
Description	<p>Configure route and next-hop properties in the event of a Bidirectional Forwarding Detection (BFD) protocol session failure event on an RSVP label-switched path (LSP). The failure event could be an existing BFD session that has gone down or a BFD session that never came up. RSVP adds back the route or next hop when the relevant BFD session comes back up.</p>
Options	<p>make-before-break—When a BFD session fails for an RSVP LSP, an attempt is made to signal a new LSP path before tearing down the old LSP path.</p> <p>teardown—When a BFD session fails for an RSVP LSP, the associated LSP path is taken down and resigaled immediately.</p> <p>teardown-timeout <i>seconds</i>—When you configure the make-before-break option, you can specify a time in seconds for the teardown-timeout option. At the end of the time specified, the associated RSVP LSP is automatically torn down and resigaled.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Failure Action for the BFD Session on an RSVP LSP on page 356

family mpls

Syntax	<pre> family mpls { all-labels; label-1; label-2; label-3; no-labels; no-label-1-exp; payload { ether-pseudowire; ip { disable; layer-3-only; port-data { source-msb; source-lsb; destination-msb; destination-lsb; } } } } </pre>
Hierarchy Level	[edit forwarding-options hash-key]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>no-label-1-exp option introduced in Junos OS Release 8.0.</p> <p>label-3 and no-labels options introduced in Junos OS Release 8.1.</p> <p>ether-pseudowire option introduced in Junos OS Release 9.1 (M320 and T Series routers only); support extended to M120 and MX Series routers in Junos OS Release 9.4.</p> <p>all-labels and payload ip disable options introduced in Junos OS Release 12.1X48R2. (PTX Series Packet Transport Routers only).</p>
Description	For aggregated Ethernet and SONET/SDH interfaces only, configure load balancing based on MPLS labels and payload. Only the IPv4 protocol is supported.
Options	<p>family mpls—(Aggregated Ethernet interfaces, aggregated SONET/SDH interfaces, and multiple equal-cost MPLS next hops only) Incorporate MPLS label and payload information into the hash key for per-flow load balancing. Only the IPv4 protocol is supported.</p> <ul style="list-style-type: none"> • all-labels—(PTX Series Packet Transport Routers only) Up to eight MPLS labels are included in the hash key to identify the uniqueness of a flow in the Packet Forwarding Engine. This is the default setting. • label-1—(M120, M320, MX Series, and T Series routers only) Include the first MPLS label into the hash key. This is used for a one-label packet for per-flow load balancing IPv4 VPLS traffic based on IP information and MPLS labels. • label-2—(M120, M320, MX Series, and T Series routers only) Include the second MPLS label into the hash key. This is used for a two-label packet for per-flow load balancing

IPv4 VPLS traffic based on IP information and MPLS labels. To use the second MPLS label in the hash key, include both the **label-1** and **label-2** statements at the **[edit forwarding-options hash-key family mpls]** hierarchy level. By default, the router provides hashing on the first and second labels. If both labels are specified, the entire first label and the first 16 bits of the second label are hashed.

- **label-3**—(M120, M320, MX Series, and T Series routers only) Include the third MPLS label into the hash key. To use the third MPLS label, include the **label-1**, **label-2**, and **label-3** statements at the **[edit forwarding-options hash-key family mpls]** hierarchy level.
- **no-labels**—Include no MPLS labels into the hash key.
- **no-label-1-exp**—(M120, M320, MX Series, and T Series routers only) The EXP bit of the first label is not used in the hash calculation to avoid reordering complications.
- **payload**—Incorporate bits from the IP payload into the hash key for per-flow load balancing Layer 2 information based on MPLS labels.
 - **disable**—(PTX Series Packet Transport Routers only) Exclude IP payload from the hash key.
 - **ether-pseudowire**—(M120, M320, MX Series, and T Series routers only) Load-balance IPv4 traffic over Layer 2 Ethernet pseudowires.
 - **ip**—Include the IP address of the IPv4 or IPv6 payload into the hash key for per-flow load balancing Layer 2 information based on MPLS labels. For the PTX Series Packet Transport Routers, this is the default setting with both Layer 3 and Layer 4 IP information included in the hash key.
 - **disable**—(PTX Series Packet Transport Routers only) Exclude IP payload from the hash key.
 - **layer-3-only**—Include only Layer 3 IP information from the IP payload data into the hash key for per-flow load balancing Layer 2 information based on MPLS labels.
 - **port-data**—(M120, M320, MX Series, and T Series routers only) Include the source and destination port field information into the hash key. By default, the most significant byte and least significant byte of the source and destination port fields are hashed. To select specific bytes to be hashed, include one or more of the **source-msb**, **source-lsb**, **destination-msb**, and **destination-lsb** options at the **[edit forwarding-options hash-key family mpls payload ip port-data]** hierarchy level. To prevent all four bytes from being hashed, include the **layer-3-only** statement at the **[edit forwarding-options hash-key family mpls payload ip]** hierarchy level.
 - **destination-lsb**—Include the least-significant byte of the destination port.
 - **destination-msb**—Include the most-significant byte of the destination port.
 - **source-lsb**—Include the least-significant byte of the source port.
 - **source-msb**—Include the most-significant byte of the source port.

Required Privilege	interface—To view this statement in the configuration.
Level	interface-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring Load Balancing Based on MPLS Labels on DPC I-Chip-Based Hardware on page 232](#)
 - [Configuring Load Balancing for Ethernet Pseudowires](#)

fast-reroute (Protocols MPLS)

Syntax	<pre>fast-reroute { (bandwidth <i>bps</i> bandwidth-percent <i>percentage</i>); (exclude [<i>group-names</i>] no-exclude); hop-limit <i>number</i>; (include-all [<i>group-names</i>] no-include-all); (include-any [<i>group-names</i>] no-include-any); }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
Description	Establish detours for the LSP so that if a node or link in the LSP fails, the traffic on the LSP can be rerouted with minimal packet loss.
Options	The statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Fast Reroute on page 226• Fast Reroute Overview on page 46• MPLS Feature Support on QFX Series and EX4600 Switches• Interprovider and Carrier-of-Carriers VPNs

fate-sharing

Syntax	<pre>fate-sharing { group <i>group-name</i> { cost <i>value</i>; from <i>address</i> <to <i>address</i>>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options], [edit routing-options], [edit routing-instances <i>routing-instance-name</i> routing-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	<p>Specify a backup path in case the primary path becomes unusable.</p> <p>You specify one or more objects with common characteristics within a group. All objects are treated as /32 host addresses. The objects can be a LAN interface, a router ID, or a point-to-point link. Sequence is insignificant.</p> <p>Changing the fate-sharing database does not affect existing established LSPs until the next CSPF reoptimization. The fate-sharing database does affect fast-reroute detour path computations.</p>
Options	<p>cost <i>value</i>—Cost assigned to the group. Range: 1 through 65,535 Default: 1</p> <p>from <i>address</i>—Address of the router or address of the LAN/NBMA interface. For example, an Ethernet network with four hosts in the same fate-sharing group would require you to list all four of the separate from addresses in the group.</p> <p>group <i>group-name</i>—Each fate-sharing group must have a name, which can have a maximum of 32 characters, including letters, numbers, periods (.), and hyphens (-). You can define up to 512 groups.</p> <p>to <i>address</i>—(Optional) Address of egress router. For point-to-point link objects, you must specify both a from and a to address.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Alternate Backup Paths Using Fate Sharing on page 62 • <i>MPLS Applications Feature Guide for Routing Devices</i>

from (Protocols MPLS)

Syntax	from <i>address</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
Description	Specify the source address to use for the LSP. The address you specify does not affect the outgoing interface used by the LSP.
Default	If you do not include this statement, the software automatically selects the loopback interface as the address.
Options	<i>address</i> —IP address.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Ingress Router Address for LSPs on page 212

gpip

Syntax	<code>gpip (ethernet hdlc ipv4 pos-scrambling-crc-16 pos-no-scrambling-crc-16 pos-scrambling-crc-32 pos-no-scrambling-crc-32 ppp);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes], [edit protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes]
Release Information	Statement introduced before Junos OS Release 7.4. pos-scrambling-crc-16 , pos-no-scrambling-crc-16 , pos-scrambling-crc-32 , and pos-no-scrambling-crc-32 options added in Junos OS Release 8.0. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the type of payload carried by the LSP. It can be any of the following: <ul style="list-style-type: none"> • ethernet—Ethernet (GPID value: 33) • hdlc—High-level Data Link Control (HDLC) (GPID value: 44) • ipv4—IP version 4 (GPID value: 0x0800) • pos-no-scrambling-crc-16—for interoperability with other vendors' equipment (GPID value: 29) • pos-no-scrambling-crc-32—for interoperability with other vendors' equipment (GPID value: 30) • pos-scrambling-crc-16—for interoperability with other vendors' equipment (GPID value: 31) • pos-scrambling-crc-32—for interoperability with other vendors' equipment (GPID value: 32) • ppp—Point-to-Point Protocol (PPP) (GPID value: 50)
Default	<code>ipv4</code>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the GPID on page 689

gre (Routing Options)

Syntax	gre;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit routing-options dynamic-tunnels <i>tunnel-name</i>]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	Enable generic routing encapsulation (GRE) type for IPv4 to automatically establish LSPs for any new PE router added to a full mesh of LSPs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MPLS-Signaled LSPs to Use GRE Tunnels on page 70

hop-limit

Syntax	<code>hop-limit <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	<p>Specify the maximum number of routers that an LSP can traverse. This limit can be applied to any of the following:</p> <ul style="list-style-type: none"> • LSPs—The configured hop limit includes the ingress and egress routers. You can specify a hop limit for an LSP and for both primary and secondary paths. • Fast reroute detour—Specify the number of additional routers a fast reroute detour can traverse relative to the protected LSP. For example, if an LSP traverses 4 routers, any detour for the LSP can be no more than 10 router hops, including the ingress and egress routers. • Link protection bypass—Specify the maximum number of routers that a link protection bypass can traverse.
Options	<p><i>number</i>—Maximum number of hops.</p> <p>Range: 2 through 255 (for an LSP or for a link protection bypass); 0 through 255 (for fast reroute)</p> <p>Default: 255 (for an LSP or for a link protection bypass); 6 (for fast reroute)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Fast Reroute on page 226 • Limiting the Number of Hops in LSPs on page 256 • Configuring the Hop Limit for Bypass LSPs on page 505

import (MPLS Traffic Engineering Database)

Syntax	<pre>import { bgp-ls-identifier <i>domain-identifier</i>; identifier <i>identifier</i>; policy <i>policy-name</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls traffic-engineering database], [edit protocols mpls traffic-engineering database]
Release Information	Statement introduced in Junos OS Release 14.2.
Description	Configure the traffic engineering database import parameters.
Options	<p>bgp-ls-identifier <i>domain-identifier</i>—BGP-TE domain identifier.</p> <p>identifier <i>identifier</i>—BGP-TE identifier. Range: 2 through 18446744073709551615</p> <p>policy <i>policy-name</i>—Name of the import policy.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• traffic-engineering on page 962

include-all (for Administrative Groups)

Syntax	<code>include-all [<i>group-names</i>];</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> admin-group],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i> admin-group],</p> <p>[edit protocols mpls label-switched-path <i>lsp-name</i> admin-group],</p> <p>[edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i> admin-group]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Require the LSP to traverse links that include all of the defined administrative groups.
Options	<i>group-names</i> —One or more names of groups defined with the admin-groups statement.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Administrative Groups for LSPs on page 240 • admin-groups on page 831

include-all (for Fast Reroute)

Syntax	<code>(include-all [<i>group-names</i>] no-include-all);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute],</p> <p>[edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.</p>
Description	<p>Control inclusion of administrative groups:</p> <ul style="list-style-type: none"> • include-all—Define the administrative groups that must all be included for fast reroute. • no-include-all—Disable administrative group inclusion.
Options	<i>group-names</i> —One or more names of groups defined with the admin-groups statement.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Fast Reroute on page 226

include-any (for Administrative Groups)

Syntax	<code>include-any [<i>group-names</i>];</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> admin-group],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i> admin-group],</code> <code>[edit protocols mpls label-switched-path <i>lsp-name</i> admin-group],</code> <code>[edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i> admin-group]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define the administrative groups to include for an LSP or for a path's primary and secondary paths.
Options	<i>group-names</i> —One or more names of groups defined with the admin-groups statement.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Administrative Groups for LSPs on page 240

include-any (for Fast Reroute)

Syntax	<code>(include-any [<i>group-names</i>] no-include-any);</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute],</code> <code>[edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.
Description	Control inclusion of administrative groups: <ul style="list-style-type: none">• include-any—Define the administrative groups to include for fast reroute.• no-include-any—Disable administrative group inclusion.
Options	<i>group-names</i> —One or more names of groups defined with the admin-groups statement.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Fast Reroute on page 226

ingress (LSP)

Syntax	<pre> ingress { bandwidth <i>bps</i>; class-of-service <i>cos-value</i>; description <i>string</i>; entropy-label; install { destination-prefix <active>; } link-protection bypass-name <i>name</i>; metric <i>metric</i>; next-hop (<i>address</i> <i>interface-name</i> <i>address/interface-name</i>); node-protection bypass-name <i>name</i> next-next-label <i>label</i>; no-install-to-address; policing { filter <i>filter-name</i>; no-auto-policing; } preference <i>preference</i>; push <i>out-label</i>; to <i>address</i>; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i>],</p> <p>[edit protocols mpls static-label-switched-path <i>lsp-name</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 10.1.</p> <p>entropy-label option introduced in Junos OS Release 14.1.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	<p>Configure an ingress LSR for a static LSP.</p> <p>The remaining statements are explained separately</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Static LSPs on page 271

install (Protocols MPLS)

Syntax	<pre>install { <i>destination-prefix</i> <active>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
Description	Associate one or more prefixes with an LSP. When the LSP is up, all the prefixes are installed as entries into the inet.3 or inet6.3 routing table.
Options	active —(Optional) Install the route into the inet.0 or inet6.0 routing table. This allows you to issue a ping or traceroute command on this address. <i>destination-prefix</i> —IPv4 or IPv6 address to associate with the LSP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Adding LSP-Related Routes to the inet.3 or inet6.3 Routing Table on page 228

ingress-policy

Syntax	<code>ingress-policy [<i>ingress-policy-names</i>];</code>
Hierarchy Level	[edit logical-system <i>logical-system-name</i> protocols ldp entropy-label], [edit logical-system <i>logical-system-name</i> protocols ldp oam], [edit protocols ldp entropy-label], [edit protocols ldp oam]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced at the [edit protocols ldp entropy-label] hierarchy level in Junos OS Release 14.1.
Description	<p>Configure an LDP ingress policy for either the entropy label or Operation, Administration, and Management (OAM).</p> <p>For OAM, configure the ingress policy to choose which forwarding equivalence classes (FECs) need to have OAM enabled. If the FEC passes through the policy or if the FEC is explicitly configured, OAM is enabled for a FEC. For FECs chosen using a policy, the BFD parameters configured under [edit protocols ldp oam bfd-liveness-detection] are applied.</p>
Options	<i>ingress-policy-names</i> —Specify the names of the ingress policies.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring OAM Ingress Policies for LDP on page 547 • Configuring the Entropy Label for LSPs on page 218

interface (Protocols MPLS)

Syntax	<pre>interface (<i>interface-name</i> all) { disable; admin-group [<i>group-names</i>]; srlg <i>srlg-name</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
Description	Enable MPLS on one or more interfaces.
Options	<p><i>interface-name</i>—Name of the interface on which to configure MPLS. To configure all interfaces, specify all. For details about specifying interfaces, see the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p> <p>srlg <i>srlg-name</i>—Name of the SRLG to associate with an interface.</p> <p>The remaining options are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Minimum MPLS Configuration on page 59• Configuring the Intermediate (Transit) and Egress Routers for Static LSPs on page 274• Example: Configuring SRLG on page 81

inter-domain

Syntax	inter-domain;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>label-switched-path-name</i>], [edit protocols mpls label-switched-path <i>label-switched-path-name</i>]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Allows the router to search for routes in the IGP database. You need to configure this statement on routers that might be unable to locate a path using intra-domain CSPF (by looking in the traffic engineering database (TED)). When you configure inter-area or inter-AS LSPs, the inter-domain statement is required.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring an LSP Across ASs on page 225 • label-switched-path on page 884

ipv6-tunneling

Syntax	ipv6-tunneling;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Allow IPv6 routes to be resolved over an MPLS network by converting LDP and RSVP routes stored in the inet.3 routing table to IPv4-mapped IPv6 addresses and then copying them into the inet6.3 routing table. This routing table can be used to resolve next hops for both inet6 and inet6-vpn routes.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks on page 71

label-switched-path (Protocols MPLS)

```
Syntax  label-switched-path lsp-name {
        disable;
        adaptive;
        admin-down;
        admin-group {
            exclude [ group-names ];
            include-all [ group-names ];
            include-any [ group-names ];
        }
        auto-bandwidth {
            adjust-interval seconds;
            adjust-threshold percentage;
            maximum-bandwidth bps;
            minimum-bandwidth bps;
            monitor-bandwidth;
        }
        bandwidth bps {
            ct0 bps;
            ct1 bps;
            ct2 bps;
            ct3 bps;
        }
        class-of-service cos-value;
        description text;
        entropy-label;
        fast-reroute {
            (bandwidth bps | bandwidth-percent percentage);
            (exclude [ group-names ] | no-exclude);
            hop-limit number;
            (include-all [ group-names ] | no-include-all);
            (include-any [ group-names ] | no-include-any);
        }
        from address;
        install {
            destination-prefix/prefix-length <active>;
        }
        inter-domain;
        ldp-tunneling;
        link-protection;
        lsp-attributes {
            encoding-type (ethernet | packet | pdh | sonet-sdh);
            gpipid (ethernet | hdlc | ipv4 | pos-scrambling-crc-16 | pos-no-scrambling-crc-16 |
                pos-scrambling-crc-32 | pos-no-scrambling-crc-32 | ppp);
            signal-bandwidth type;
            switching-type (fiber | lambda | psc-1 | tdm);
        }
        metric metric;
        no-cspf;
        no-decrement-ttl;
        node-link-protection;
        optimize-timer seconds;
        p2mp lsp-name;
```

```

policing {
    filter filter-name;
    no-auto-policing;
}
preference preference;
primary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    priority setup-priority reservation-priority;
    (record | no-record);
    select (manual | unconditional);
    standby;
}
priority setup-priority reservation-priority;
(random | least-fill | most-fill);
(record | no-record);
retry-limit number;
retry-timer seconds;
revert-timer seconds;
secondary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    priority setup-priority reservation-priority;
    (record | no-record);
}

```

```
    select (manual | unconditional);
    standby;
  }
  soft-preemption;
  standby;
  to address;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
}
```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure an LSP to use in dynamic MPLS. When configuring an LSP, you must specify the address of the egress router in the to statement. All remaining statements are optional.
Options	<p>lsp-name—Name that identifies the LSP. The name can be up to 64 characters and can contain letters, digits, periods, and hyphens. To include other characters, enclose the name in quotation marks. The name must be unique within the ingress router.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Minimum MPLS Configuration on page 59• Configuring the Ingress and Egress Router Addresses for LSPs on page 212• Configuring Primary and Secondary LSPs on page 214

label-switched-path-template (Container LSP)

Syntax	label-switched-path-template { (default-template lsp-template-name); }
Hierarchy Level	[edit protocols mpls container-label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in Junos OS Release 14.2.
Description	Specify the LSP template. An LSP template is used as the basis for other dynamically generated LSPs.
Options	<p>default-template—Specify that the default LSP template be used for the dynamically generated LSPs.</p> <p>lsp-template-name—Specify the name of an LSP to be used as a template for the dynamically generated LSPs.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • container-label-switched-path on page 848

ldp-tunneling

Syntax	ldp-tunneling;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable the LSP to be used for LDP tunneling.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Enabling LDP over RSVP-Established LSPs on page 629

least-fill

See [random](#)

link-protection (Dynamic LSPs)

Syntax	link-protection;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.
Description	<p>Enable link protection on the specified LSP, which helps to ensure that traffic sent over a specific interface to a neighboring router can continue to reach the router if that interface fails. For point-to-multipoint LSPs, including this statement extends link protection to all of the paths used by the LSP.</p> <p>To fully enable link protection, you must also include the link-protection statement at the [edit protocols rsvp interface <i>interface-name</i>] or [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>] hierarchy level.</p>
Default	Link protection is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Link Protection for Point-to-Multipoint LSPs on page 304• Configuring Node Protection or Link Protection for LSPs on page 509• link-protection (RSVP) on page 992

link-protection (Static LSPs)

Syntax	link-protection bypass-name <i>name</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress], [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>], [edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress], [edit protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	Enable link protection on the specified static LSP. Link protection helps to ensure that traffic sent over a specific interface to a neighboring router can continue to reach the router if that interface fails.
Default	Link protection is disabled.
Options	bypass-name <i>name</i> —Bypass LSP name.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Static LSPs on page 271 • <i>Example: Configuring Point-to-Multipoint LSPs with Static Routes</i>

load-balance-label-capability

Syntax	load-balance-label-capability;
Hierarchy Level	[edit forwarding-options]
Release Information	Statement introduced in Junos OS Release 14.1.
Description	Enables the router to push and pop the load balancing label and causes LDP and RSVP to advertise the entropy label TLV to neighboring routers.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Entropy Label for LSPs on page 218

log-updown (Protocols MPLS)

Syntax	<pre>log-updown { no-trap { mpls-lsp-traps; rfc3812-traps; } (syslog no-syslog); trap; trap-path-down; trap-path-up; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4. The mpls-lsp-traps and rfc-3812-traps options added in Junos OS Release 9.0. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
Description	Log a message or send an SNMP trap whenever an LSP makes a transition from up to down, or vice versa, and whenever an LSP switches from one active path to another. Only the ingress router performs these operations.
Default	There is no default behavior for this statement. If you do not specify the options, the configuration cannot be committed.
Options	no-syslog —Do not log a message to the system log file. no-trap —Do not send an SNMP trap. syslog —Log a message to the system log file. trap —Send an SNMP trap. trap-path-down —Send an SNMP trap when an LSP path goes down. trap-path-up —Send an SNMP trap when an LSP path comes up. The no-trap statement is explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring System Log Messages and SNMP Traps for LSPs on page 344• <i>Network Management Administration Guide for Routing Devices</i>• no-trap on page 908• traceoptions (Protocols MPLS) on page 954

loss (querier)

Syntax	<pre> loss { traffic-class tc-value { average-sample-size <i>sample size</i>; loss-threshold <i>loss threshold value</i>; loss-threshold-window <i>number of samples for loss threshold</i>; measurement-quantity <i>bytes packets</i>; query-interval <i>milliseconds</i>; } } </pre>
Hierarchy Level	<p>[edit protocols mpls oam performance-monitoring querier], [edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier], [edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring querier], [edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring querier]</p>
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>Configure loss measurement options.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Pro-Active Loss and Delay Measurements on page 448 • On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview on page 50 • performance-monitoring (Protocols MPLS) on page 922

loss (responder)

Syntax	<pre>loss { min-query-interval <i>milliseconds</i>; }</pre>
Hierarchy Level	[edit protocols mpls oam performance-monitoring responder], [edit protocols mpls label-switched-path <i>lsp-name</i> oam performance-monitoring responder], [edit protocols mpls label-switched-path <i>lsp-name</i> primary <i>path-name</i> oam performance-monitoring responder], [edit protocols mpls label-switched-path <i>lsp-name</i> secondary <i>path-name</i> oam performance-monitoring responder]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	Configure loss measurement options.
Options	<p>min-query-interval <i>milliseconds</i>—(Optional) Specify the minimum query interval that the responder supports. If the minimum query interval of the responder is greater than the query interval configured at the querier, the effective message query rate is the minimum query interval configured for the responder.</p> <p>Default: 10 seconds</p> <p>Range: 1000 through 4294967295 milliseconds</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Pro-Active Loss and Delay Measurements on page 448• On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview on page 50• performance-monitoring (Protocols MPLS) on page 922

loss-delay (querier)

Syntax	<pre> loss-delay { traffic-class tc-value { average-sample-size sample size; loss-threshold loss threshold value; loss-threshold-window number of samples for loss threshold; measurement-quantity bytes packets; padding-size size; query-interval milliseconds; rtt-delay-threshold rtt threshold value; twcd-delay-threshold twcd threshold value; } } </pre>
Hierarchy Level	<p>[edit protocols mpls oam performance-monitoring querier], [edit protocols mpls label-switched-path lsp-name oam performance-monitoring querier], [edit protocols mpls label-switched-path lsp-name primary path-name oam performance-monitoring querier], [edit protocols mpls label-switched-path lsp-name secondary path-name oam performance-monitoring querier]</p>
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>Configure combined loss-delay measurement options.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Pro-Active Loss and Delay Measurements on page 448 • On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview on page 50 • performance-monitoring (Protocols MPLS) on page 922

lsp-attributes

Syntax	<pre>lsp-attributes { encoding-type (ethernet packet pdh sonet-sdh); gpid (ethernet hdlc ipv4 pos-scrambling-crc-16 pos-no-scrambling-crc-16 pos-scrambling-crc-32 pos-no-scrambling-crc-32 ppp); signal-bandwidth type; switching-type (fiber lambda psc-1 tdm); }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. pos-scrambling-crc-16 , pos-no-scrambling-crc-16 , pos-scrambling-crc-32 , and pos-no-scrambling-crc-32 options added in Junos OS Release 8.0. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Define the parameters signaled during LSP setup. These usually determine the nature of the resource (label) allocated for the LSP. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring MPLS LSPs for GMPLS on page 688

maximum-bandwidth (Protocols MPLS)

Syntax	maximum-bandwidth <i>bps</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Specify the maximum amount of bandwidth in bits per second (bps).
Options	<i>bps</i> —Maximum amount of bandwidth.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 259

maximum-labels

Syntax	<code>maximum-labels <i>maximum-labels</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	<p>On the logical interface, specify the maximum number of MPLS labels upon which MPLS can operate.</p> <p>You can configure this statement on the following routers:</p> <ul style="list-style-type: none"> • MX Series 3D Universal Edge Router • M120 Multiservice Edge Router • M320 Multiservice Edge Router with Enhanced III FPCs • M7i Multiservice Edge Router and M10i Multiservice Edge Router with Enhanced Compact Forwarding Engine Board (CFEB-E) • T640, T1600, T4000, TX Matrix, and TX Matrix Plus routers with Enhanced Scaling FPC1, Enhanced Scaling FP2, Enhanced Scaling FPC3, and Enhanced Scaling FPC4
Options	<p><i>maximum-labels</i>—Maximum number of labels.</p> <p>Range: 3 through 5</p> <p>Default: 3</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Maximum Number of MPLS Labels on page 334 • <i>Junos OS VPNs Library for Routing Devices</i>

minimum-bandwidth-adjust-interval

Syntax	<code>minimum-bandwidth-adjust-interval <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Specify the duration (in seconds) for which minimum bandwidth is frozen.
Options	<i>seconds</i> —Minimum bandwidth reallocation interval, in seconds. Range: 300 through 31,536,000 seconds.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 259

minimum-bandwidth-adjust-threshold-change

Syntax	<code>minimum-bandwidth-adjust-threshold-change <i>percentage</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Specify the percentage change in maximum average bandwidth to freeze the minimum bandwidth.
Options	<i>percentage</i> —Percentage change in maximum average bandwidth. Range: Range: 0 through 100 percent.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 259

minimum-bandwidth-adjust-threshold-value

Syntax	minimum-bandwidth-adjust-threshold-value <i>bps</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Specify the value in bits per second (bps) to freeze the minimum bandwidth if the maximum average bandwidth falls below this value.
Options	<i>bps</i> —Threshold value for minimum bandwidth if the maximum average bandwidth falls below the specified value.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 259

metric (Protocols MPLS)

Syntax	<code>metric <i>metric</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
Description	Compare against another LSP or against an IGP route. To disable dynamic metric tracking, assign a fixed metric value to an LSP. If no metric is assigned, the LSP metric is dynamic and automatically tracks underlying IGP metrics.
Options	<i>metric</i> —LSP metric value. Default: No metric assigned (dynamic) Range: 1 through 16,777,215
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Static LSP Metrics on page 231

minimum-bandwidth

Syntax	<code>minimum-bandwidth <i>bps</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Set the minimum bandwidth in bps for an LSP with automatic bandwidth allocation enabled.
Options	<i>bps</i> —Minimum bandwidth for the LSP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 259

monitor-bandwidth

Syntax	monitor-bandwidth;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Do not automatically adjust bandwidth allocation. However, the maximum average bandwidth utilization is monitored on the LSP, and the information is recorded in the MPLS statistics file.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Passive Bandwidth Utilization Monitoring on page 263

most-fill

See [random](#)

mpls (Protocols)

Syntax	mpls { ... }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit protocols]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable MPLS on the router.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Minimum MPLS Configuration on page 59

mpls-tp-mode

Syntax	mpls-tp-mode;
Hierarchy Level	[edit protocols mpls label-switched-path <i>lsp-name</i> oam], [edit protocols mpls oam]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	<p>Enable GAL or G-Ach OAM operation without IP encapsulation on a label-switched path (LSP).</p> <p>Include this statement at the [edit protocols mpls oam] hierarchy level to enable GAL or G-Ach OAM operation without IP encapsulation on all LSPs in the MPLS network. Include this statement at the [edit protocols mpls label-switched-path <i>lsp-name</i> oam] hierarchy level to enable GAL and G-Ach OAM operation without IP encapsulation on a specific LSP.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring the MPLS Transport Profile for OAM on page 136

mtu-signaling

Syntax	mtu-signaling;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls path-mtu rsvp], [edit protocols mpls path-mtu rsvp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable MTU signaling in RSVP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling MTU Signaling in RSVP on page 485

next-hop (Protocols MPLS)

Syntax	<code>next-hop (address interface-name address/interface-name);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> bypass],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>],</p> <p>[edit protocols mpls static-label-switched-path <i>lsp-name</i> bypass],</p> <p>[edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress],</p> <p>[edit protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Location of the next hop to the destination, specified as the IPv4 address of the next hop, the interface name (for point-to-point interfaces only), or the address/interface-name to specify an IP address on an operational interface.
Options	<p>address—IPv4 address of the next-hop router.</p> <p>interface-name—IP address of the outgoing interface. It must be a point-to-point interface. The name can be a simple or fully qualified domain name.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Ingress Router for Static LSPs on page 271

no-bfd-triggered-local-repair

Syntax	no-bfd-triggered-local-repair;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
Release Information	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Disable Bidirectional Forwarding Detection (BFD) sessions to trigger fast reroute (FRR) using MPLS-FRR and loop-free alternates (LFAs). When this statement is configured, no BFD-triggered local repair is supported. However, logical interface down-based local repair is in force.</p> <p>When using this statement to disable local repair, you also must restart routing to ensure proper behavior. To restart routing, include the graceful-restart command for the interior gateway protocol (IGP) used in your configuration. For example, if your IGP is OSPF, include the graceful-restart statement at the [edit protocols ospf] hierarchy level.</p>
Default	BFD-triggered local repair is the default behavior. The loss of a neighbor results in BFD local repair for all next hops that derive themselves from the base next hop with which the BFD session is established.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• BFD-Triggered Local Repair for Rapid Convergence on page 357• <i>graceful-restart (Enabling Globally)</i>

no-cspf

Syntax	no-cspf;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	<p>Disable constrained-path LSP computation.</p> <p>An explicit-path LSP is completely configured through operator action. Once configured, it is initiated only along the explicitly specified path.</p> <p>A constrained-path LSP relies on an ingress router to compute the complete path. The ingress router takes into account the following information during the computation:</p> <ul style="list-style-type: none"> • Interior gateway protocol (IGP) topology database • Link utilization information from extensions in the IGP link-state database • Administrative group information from extensions in the IGP link-state database • LSP requirements, including bandwidth, hop count, and administrative group <p>Constrained-path LSPs can generally avoid link failures and congested links. They also permit recomputation (therefore, a new path) during topology changes or unsuccessful setup.</p>
Default	Constrained-path LSP computation enabled.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Disabling Constrained-Path LSP Computation on page 239 • Configuring Explicit-Path LSPs on page 278

no-decrement-ttl

Syntax	no-decrement-ttl;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Disable normal time-to-live (TTL) decrementing, which decrements the TTL field in the IP header by 1. This statement decrements the IP TTL by 1 before encapsulating the IP packet within an MPLS packet. When the penultimate router pops off the top label, it does not use the standard write-back procedure of writing the MPLS TTL into the IP TTL field. Therefore, the IP packet is decremented by 1. The ultimate router then decrements the packet by one more for a total cloud appearance of 2, thus hiding the network topology.
Default	Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Disabling Normal TTL Decrementing on page 236• no-propagate-ttl on page 907

no-install-to-address

Syntax	no-install-to-address;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress]
Release Information	Statement introduced in Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Prevent the egress router address configured using the to statement from being installed into the inet.3 and inet.0 routing tables.
Default	The egress router address for an LSP is installed into the inet.3 and inet.0 routing tables.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Preventing the Addition of Egress Router Addresses to Routing Tables on page 213 • to on page 953

no-load-balance-label-capability

Syntax	no-load-balance-label-capability;
Hierarchy Level	[edit forwarding-options]
Release Information	Statement introduced in Junos OS Release 14.1.
Description	Disables advertisement of entropy label capability in LDP and RSVP.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • entropy-label on page 860 • Configuring the Entropy Label for LSPs on page 218

no-mcast-replication

Syntax	no-mcast-replication;
Hierarchy Level	[edit chassis fpc <i>slot-number</i> pic <i>pic-number</i>], [edit chassis lcc <i>number</i> fpc <i>slot-number</i> pic <i>pic-number</i>]
Release Information	Statement introduced in Junos OS Release 11.3.
Description	For point-to-multipoint LSPs configured on T Series routers, protect the Packet Forwarding Engine (PFE) from bandwidth saturation. When a PFE does not need to replicate traffic, the PFE's bandwidth is less likely to become saturated. When you include the no-mcast-replication statement, the PFE is forced to be a leaf node in the binary tree. Leaf nodes, unlike branch nodes, do not replicate traffic in the process of forwarding traffic. Because leaf nodes have no children, they do not need to replicate traffic, and thus are less likely to become saturated with traffic.
Default	If you omit the no-mcast-replication statement, the PFE can become a branch node or a leaf node. When the PFE becomes a branch node, the PFE must replicate traffic.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Point-to-Multipoint LSPs Overview on page 281

no-propagate-ttl

Syntax	no-propagate-ttl;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	<p>Disable normal time-to-live (TTL) decrementing. You configure this statement once per router, and it affects all RSVP-signaled or LDP-signaled LSPs. When this router acts as an ingress router for an LSP, it pushes an MPLS header with a TTL value of 255, regardless of the IP packet TTL. When the router acts as the penultimate router, it pops the MPLS header without writing the MPLS TTL into the IP packet.</p> <p>When you add the no-propagate-ttl statement to the configuration or delete it from the configuration, the effect takes place immediately. There is no need to clear existing RSVP LSPs or LDP sessions.</p>
Default	Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Disabling Normal TTL Decrementing on page 236 • <i>Example: Disabling Normal TTL Decrementing in a VRF Routing Instance (on Layer 3 VPNs Feature Guide for Routing Devices or in the Junos VPNs Configuration Guide)</i> • no-decrement-ttl on page 904

no-transit-statistics

Syntax	no-transit-statistics;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls statistics], [edit protocols mpls statistics]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Disables the collection of MPLS statistics for LSPs transiting the router. You cannot configure this statement and the transit-statistics-polling statement at the same time.
Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring MPLS to Gather Statistics on page 342 • statistics on page 947

no-trap

Syntax	<pre>no-trap { mpls-lsp-traps; rfc-3812-traps; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls log-updown], [edit protocols mpls log-updown]
Release Information	Statement introduced before Junos OS Release 7.4. The mpls-lsp-traps and rfc-3812-traps options added in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Prevent the transmission of SNMP traps.
Options	mpls-lsp-traps —Block the MPLS LSP traps defined in the rfc-3812-traps , but allows the rfc3812.mib traps. rfc-3812-traps —Block the traps defined in the rfc3812.mib , but allows the MPLS LSP traps defined in the jnx-mpls.mib .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring System Log Messages and SNMP Traps for LSPs on page 344• <i>Network Management Administration Guide for Routing Devices</i>• traceoptions (Protocols MPLS) on page 954

node-protection (Static LSP)

Syntax	<code>node-protection bypass-name <i>name</i> next-next-label <i>label</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>],</p> <p>[edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress],</p> <p>[edit protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>]</p>
Release Information	Statement introduced in JUNOS Release 10.1.
Description	Enable node protection on the specified static bypass LSP. Node protection ensures that traffic from an LSP traversing a neighboring router can continue to reach its destination even if the neighboring router fails.
Default	Node protection is disabled.
Options	<p><code>bypass-name <i>name</i></code>—Bypass LSP name.</p> <p><code>next-next-label <i>label</i></code>—Bypass LSP name.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Static LSPs on page 271

normalization

Syntax	<pre>normalization { failover-normalization; no-incremental-normalize; normalization-retry-duration <i>seconds</i>; normalization-retry-limits <i>number</i>; normalize-interval <i>seconds</i>; }</pre>
Hierarchy Level	[edit protocols mpls container-label-switched-path <i>lsp-name</i> splitting-merging]
Release Information	Statement introduced in Junos OS Release 14.2.
Description	Perform normalization.
Options	<p>failover-normalization—Enable the ingress router to pro-actively normalize or re-distribute traffic when a link or node failure happens on a member LSP. A member LSP can go down between two scheduled normalization events because of a link-failure or pre-emption.</p> <p>Default: Disabled</p> <p>no-incremental-normalize—Disables automatic switchover by the ingress router to a new instance of the container LSP until the desired demand is satisfied, although the given number of LSPs can be successfully signaled such that the new aggregate bandwidth value exceeds the old aggregate bandwidth value.</p> <p>Default: False (disabled)</p> <p>normalization-retry-duration <i>seconds</i>—Specifies the duration before which the ingress router performs a normalization reattempt when the previous normalization has not been successful. Normalization is done until a sufficient number of LSPs come up with an aggregate bandwidth that is more than the current aggregate or desired bandwidth.</p> <p>Default: 30 seconds</p> <p>normalization-retry-limits <i>number</i>—Specifies the maximum number of times the ingress router performs normalization reattempts until a sufficient number of LSPs come up successfully with new bandwidth values.</p> <p>Default: 1</p> <p>normalize-interval <i>seconds</i>—Specifies the duration between two normalization events.</p> <p>Range: 21600 seconds through 6 hours</p> <p>Default: 21600 seconds</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• splitting-merging on page 941

oam (Protocols MPLS)

```
Syntax  oam {
    bfd-liveness-detection{
        failure-action teardown;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        minimum-transmit-interval milliseconds;
        multiplier detection-time-multiplier;
    }
    lsp-ping-interval seconds;
    mpls-tp-mode;
    performance-monitoring {
        querier {
            loss {
                traffic-class tc-value {
                    query-interval milliseconds;
                    measurement-quantity bytes|packets;
                    average-sample-size sample size;
                    loss-threshold loss threshold value;
                    loss-threshold-window number of samples for loss threshold;
                }
            }
            delay {
                traffic-class tc-value {
                    query-interval milliseconds;
                    padding-size size;
                    average-sample-size sample size;
                    rtt-delay-threshold rtt threshold value;
                    twcd-delay-threshold twcd threshold value;
                }
            }
            loss-delay {
                traffic-class tc-value {
                    query-interval milliseconds;
                    measurement-quantity bytes|packets;
                    padding-size size;
                    average-sample-size sample size;
                    loss-threshold loss threshold value;
                    loss-threshold-window number of samples for loss threshold;
                    rtt-delay-threshold rtt threshold value;
                    twcd-delay-threshold twcd threshold value;
                }
            }
        }
    }
    responder {
        loss {
            min-query-interval milliseconds;
        }
        delay {
            min-query-interval milliseconds;
        }
    }
}
```

	}
Hierarchy Level	[edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>] [edit protocols mpls label-switched-path <i>lsp-name</i> primary <i>path-name</i>]
Release Information	Statement introduced in Junos OS Release 7.6. lsp-ping-interval option introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric. performance-monitoring configuration statement introduced in Junos OS Release 15.1.
Description	Enable Operation, Administration, and Maintenance (OAM) for RSVP-signaled LSPs.
Options	lsp-ping-interval <i>seconds</i> —Specify the duration of the LSP ping interval in seconds. To issue a ping on an RSVP-signaled LSP, use the ping mpls rsvp command. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BFD for MPLS IPv4 LSPs on page 354

optimize-adaptive-teardown

Syntax	optimize-adaptive-teardown { p2p: }
Hierarchy Level	[edit protocols mpls]
Release Information	Statement introduced in Junos OS Release 15.1R1.
Description	Make use of a new feedback mechanism from TAG library which relies on RPD infrastructure to decide when all the routes using the old LSP instance have fully shifted to the new LSP instance after MBB switchover. When this statement is configured, the optimize-hold-dead-delay statement, which delays the teardown of the old LSP instance after MBB switchover, is ignored.
Options	p2p —This is the only option. Only point-to-point LSPs configured in the system will be affected.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Optimizing Signaled LSPs on page 251• Achieving a Make-Before-Break, Hitless Switchover for LSPs on page 246

optimize-aggressive

Syntax	optimize-aggressive;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	If enabled, the LSP reoptimization is based solely on the IGP metric. The reoptimization process ignores the available bandwidth ratio calculations, the least-fill 10 percent congestion improvement rule, and the hop-counts rule. This statement makes reoptimization more aggressive than the default.
Default	Aggressive optimization is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Optimizing Signaled LSPs on page 251

optimize-hold-dead-delay

Syntax	<code>optimized-hold-dead-delay seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switch-path <i>lsp-name</i>], [edit protocols mpls], [edit protocols mpls label-switch-path <i>lsp-name</i>]
Description	Allows you to specify the amount of time to delay the tear down of old paths after the router has switched traffic to new optimized paths. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The specified delay helps to ensure that old paths are not torn down before all routes have been switched over to the new optimized paths. This delay timer starts when the timer specified by the optimize-switchover-dealy statement has elapsed.
Options	seconds —Configure the time in seconds to wait before tearing down the old paths that were in use prior to the last LSP optimization. Default: 60 seconds Range: 0 through 65,535 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Optimizing Signaled LSPs on page 251• optimize-switchover-delay on page 915• optimize-timer on page 916

optimize-switchover-delay

Syntax	<code>optimize-switchover-delay <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced in Junos OS Release 11.1R1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Delays the switch over of LSPs to newly optimized paths. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The specified delay helps to ensure that the new optimized paths have been established before traffic is switched over from the old paths.
Options	<p><i>seconds</i>—Configure the time in seconds to wait before switching LSPs to newly optimized paths.</p> <p>Default: 1 second</p> <p>Range: 1 through 900 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Optimizing Signaled LSPs on page 251 • optimize-hold-dead-delay on page 914 • optimize-timer on page 916

optimize-timer (Protocols MPLS)

Syntax	<code>optimize-timer <i>seconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.</p>
Description	<p>Enable periodic reoptimization of an LSP that is already set up. If topology changes occur, an existing path might become suboptimal, and a subsequent recomputation might be able to determine a better path. This feature is useful only on LSPs for which constrained-path computation is enabled; that is, for which the no-cspf statement is not configured. Also, you only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers).</p> <p>To avoid extensive resource consumption that might result because of frequent path recomputations, or to avoid destabilizing the network as a result of constantly changing LSPs, we recommend that you either leave the timer value sufficiently large or disable the timer value.</p>
Default	The optimize timer is disabled.
Options	<p>seconds—Length of the optimize timer, in seconds.</p> <p>Range: 0 through 65,535 seconds</p> <p>Default: 0 seconds (the optimize timer is disabled)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Optimizing Signaled LSPs on page 251

p2mp (Protocols MPLS)

Syntax	<code>p2mp p2mp-lsp-name;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
Description	Specify an LSP as either a point-to-multipoint LSP or as a branch LSP of a point-to-multipoint LSP by specifying the point-to-multipoint LSP path name.
Options	<i>p2mp-lsp-name</i> —Name of the point-to-multipoint LSP path that identifies the sequence of nodes that form the point-to-multipoint LSP. The name can contain up to 32 characters and can include letters, digits, periods, and hyphens. To include other characters or use a longer name, enclose the name in quotation marks. The name must be unique within the ingress router.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Primary Point-to-Multipoint LSP on page 283

p2mp-lsp-next-hop

Syntax	<pre>p2mp-lsp-next-hop { metric <i>metric</i>; preference <i>preference</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static route <i>destination-prefix</i>], [edit logical-systems <i>logical-system-name</i> routing-options static route <i>destination-prefix</i>], [edit routing-instances <i>routing-instance-name</i> routing-options static route <i>destination-prefix</i>]. [edit routing-options static route <i>destination-prefix</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	<p>Specify a point-to-multipoint LSP as the next hop for a static route, and configure an independent metric or preference on that next-hop LSP.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Static Unicast Routes for Point-to-Multipoint LSPs on page 277• Example: Configuring a Collection of Paths to Create an RSVP-Signaled Point-to-Multipoint LSP on page 285• Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on Logical Systems

path (Protocols MPLS)

Syntax	<pre>path <i>path-name</i> { (<i>address</i> <i>hostname</i>) <strict loose>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
Description	<p>Create a named path and optionally specify the sequence of explicit routers that form the path.</p> <p>You must include this statement when configuring explicit LSPs.</p>
Options	<p>address—IP address of each transit router in the LSP. You must specify the address or hostname of each transit router, although you do not need to list each transit router if its type is loose. As an option, you can include the ingress and egress routers in the path. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path until reaching the egress router (optional) or the router immediately before the egress router.</p> <p>Default: If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.</p> <p>hostname—See address.</p> <p>Default: If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.</p> <p>loose—(Optional) Indicate that the next address in the path statement is a loose link. This means that the LSP can traverse through other routers before reaching this router.</p> <p>Default: strict</p> <p>path-name—Name that identifies the sequence of nodes that form an LSP. The name can contain up to 32 characters and can include letters, digits, periods, and hyphens. To include other characters or use a longer name, enclose the name in quotation marks. The name must be unique within the ingress router.</p> <p>strict—(Optional) Indicate that the LSP must go to the next address specified in the path statement without traversing other nodes. This is the default.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Creating Named Paths on page 60

path-mtu

Syntax	<pre>path-mtu { allow-fragmentation; rsvp { mtu-signaling; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Configure MTU options for MPLS paths, including packet fragmentation and MTU signaling. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MTU Signaling in RSVP on page 484

per-prefix-label

Syntax	<code>per-prefix-label;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family inet labeled-unicast], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit protocols bgp family inet labeled-unicast], [edit protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit routing-instances <i>instance-name</i> logical-systems <i>logical-system-name</i> protocols bgp family inet labeled-unicast], [edit routing-instances <i>instance-name</i> logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit routing-instances <i>instance-name</i> protocols bgp family inet labeled-unicast], [edit routing-instances <i>instance-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1x48 for PTX Series Packet Transport Routers.</p> <p>Statement introduced in Junos OS Release 12.3 for M Series, T Series, and MX Series routers.</p>
Description	<p>Allocate a unique label for each prefix. The per-prefix-label statement helps minimize packet loss in most deployments.</p> <p>Although allocating a label for each prefix is not generally ideal for scaling, it is assumed that a small number of labels are used for BGP labeled-unicast. When labeled BGP is used to set up transport label-switched paths (LSPs), the common case is that each prefix has a unique next hop. Thus, the use of per-prefix labels does not have an adverse scaling impact. On the contrary, the use of per-prefix labels reduces churn in the network when multipath load balancing is enabled for IPv4 labeled-unicast, and a subset of the paths are withdrawn for some reason.</p> <p>The advantage of per-prefix labeling is that the advertised upstream label is more stable during network changes. That is, if the downstream label changes, the advertised upstream label remains the same under most scenarios. This way, the upstream router is isolated from the downstream network change, and the overall network is more stable. The greater stability of the advertised upstream label helps to reduce traffic loss during many different network change scenarios.</p>
Default	By default, label allocation is per next-hop router.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • MPLS Label Allocation on page 26

performance-monitoring (Protocols MPLS)

Syntax	<pre> performance-monitoring { querier { delay { traffic-class tc-value { average-sample-size sample size; padding-size size; query-interval milliseconds; rtt-delay-threshold rtt threshold value; twcd-delay-threshold twcd threshold value; } } loss { traffic-class tc-value { average-sample-size sample size; loss-threshold loss threshold value; loss-threshold-window number of samples for loss threshold; measurement-quantity bytes packets; query-interval milliseconds; } } loss-delay { traffic-class tc-value { average-sample-size sample size; loss-threshold loss threshold value; loss-threshold-window number of samples for loss threshold; measurement-quantity bytes packets; padding-size size; query-interval milliseconds; rtt-delay-threshold rtt threshold value; twcd-delay-threshold twcd threshold value; } } } responder { delay { min-query-interval milliseconds; } loss { min-query-interval milliseconds; } } } </pre>
Hierarchy Level	<pre> [edit protocols mpls oam], [edit protocols mpls label-switched-path lsp-name oam], [edit protocols mpls label-switched-path lsp-name primary path-name oam], [edit protocols mpls label-switched-path lsp-name secondary path-name oam] </pre>
Release Information	Statement introduced in Junos OS Release 15.1.
Description	Configure performance monitoring options.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

-

policing (Protocols MPLS)

Syntax `policing {
 filter filter-name;
 no-auto-policing;
}`

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls label-switched-path *lsp-name*],
[edit logical-systems *logical-system-name* protocols mpls
 static-label-switched-path *lsp-name* ingress],
[edit protocols mpls label-switched-path *lsp-name*],
[edit protocols mpls static-label-switched-path *lsp-name* ingress]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Specify the policing filter for the LSP.

Options `filter filter-name`—Specify the name of the policing filter.
`no-auto-policing`—Disable automatic policing on this LSP.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring Policers for LSPs on page 347](#)
- [auto-policing on page 840](#)

pop

Syntax	pop;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>], [edit protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Remove the label from the top of the label stack. If there is another label in the stack, that label becomes the label at the top of the label stack. Otherwise, the packet is forwarded as a native protocol packet (typically, as an IP packet).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Intermediate (Transit) and Egress Routers for Static LSPs on page 274• swap on page 948

preference (Protocols MPLS)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	<p>Preference for the route.</p> <p>You can optionally configure multiple LSPs between the same pair of ingress and egress routers. This is useful for balancing the load among the LSPs because all LSPs, by default, have the same preference level. To prefer one LSP over another, set different preference levels for individual LSPs. The LSP with the lowest preference value is used. The default preference for LSPs is lower (more preferred) than all learned routes except direct interface routes.</p>
Options	<p>preference—Preference to assign to the route. A route with a lower preference value is preferred.</p> <p>Range: 1 through 255</p> <p>Default: 5 for static MPLS LSPs, 7 for RSVP MPLS LSPs, 9 for LDP MPLS LSPs</p>
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Preference Values for LSPs on page 243 • Configuring the Ingress Router for Static LSPs on page 271 • Configuring the Intermediate (Transit) and Egress Routers for Static LSPs on page 274

primary (Protocols MPLS)

Syntax	<pre> primary <i>path-name</i> { adaptive; admin-group { exclude [<i>group-names</i>]; include-all [<i>group-names</i>]; include-any [<i>group-names</i>]; } bandwidth <i>bps</i>; class-of-service <i>cos-value</i>; hop-limit <i>number</i>; no-cspf; no-decrement-ttl; optimize-timer <i>seconds</i>; preference <i>preference</i>; priority <i>setup-priority reservation-priority</i>; (record no-record); select (manual unconditional); standby; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	<p>Specify the primary path to use for an LSP. You can configure only one primary path.</p> <p>You can optionally specify preference, CoS, and bandwidth values for the primary path, which override any equivalent values that you configure for the LSP (at the [edit mpls label-switched-path <i>lsp-name</i>] hierarchy level).</p>
Options	<p><i>path-name</i>—Name of a path that you created with the path statement.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Primary and Secondary LSPs on page 214

priority (Protocols MPLS)

Syntax	<code>priority setup-priority reservation-priority;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the setup priority and reservation priority for an LSP. If insufficient link bandwidth is available during session establishment, the setup priority is compared with other setup priorities for established sessions on the link to determine whether some of them should be preempted to accommodate the new session. Sessions with lower hold priorities are preempted.
Options	<p>reservation-priority—Reservation priority, used to keep a reservation after it has been set up. A smaller number has a higher priority. The priority must be greater than or equal to the setup priority to prevent preemption loops.</p> <p>Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p>Default: 0 (Once the session is set up, no other session can preempt it.)</p> <p>setup-priority—Setup priority.</p> <p>Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p>Default: 7 (The session cannot preempt any existing sessions.)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Priority and Preemption for LSPs on page 250

protection-revert-time

Syntax	<code>protection-revert-time <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static], [edit protocols mpls interface <i>interface-name</i> static]
Release Information	Statement introduced in Junos OS Release 10.1.
Description	<p>Specify the amount of time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path.</p> <p>If you have configured a value of 0 seconds for the protection-revert-time statement and traffic is switched to the bypass path, the traffic remains on that path indefinitely. It is never switched back to the original path unless the bypass path is down or you intervene.</p>
Options	<p><i>seconds</i>—Time in seconds.</p> <p>Range: 0 through 65,535 seconds</p> <p>Default: 5 seconds</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Static LSPs on page 271

push

Syntax	<code>push out-label;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> bypass], [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress], [edit protocols mpls static-label-switched-path <i>lsp-name</i> bypass], [edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Add a new label to the top of the label stack. This statement is used to configure static LSPs at ingress routers and to configure bypass LSPs for static LSPs.
Options	out-label —Manually assigned outgoing label value. Range: 0 through 1,048,575.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • pop on page 924 • swap on page 948 • Configuring the Ingress Router for Static LSPs on page 271

random

Syntax	(random least-fill most-fill);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Configure the preferred path when several equal-cost candidate paths to a destination exist, and prefer the path with the highest available bandwidth (with the largest minimum available bandwidth ratio). The available bandwidth ratio of a link is the available bandwidth on a link divided by the maximum reservable bandwidth on the link.</p> <ul style="list-style-type: none">• least-fill—Prefer the path with the most available bandwidth (with the largest available bandwidth ratio).• most-fill—Prefer the path with the least available bandwidth (with the minimum available bandwidth ratio). The minimum available bandwidth ratio of a path is the smallest available bandwidth ratio belonging to any of the links in the path.• random—Choose the path at random.
Default	random
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring CSPF Tie Breaking on page 232

record

Syntax	(record no-record);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify whether an LSP should actively record the routes in the path. Recording routes requires that all transit routers support the RSVP Record Route object. Recording routes can be useful for diagnostics and loop detection.
Default	Record routes.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Disabling Path Route Recording by LSPs on page 244

retry-limit

Syntax	<code>retry-limit <i>number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>],
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Maximum number of times the ingress router tries to establish the primary path. This counter is reset each time a primary path is created successfully. When the limit is exceeded, no more connection attempts are made. Intervention is then required to restart the connection.
Options	<i>number</i> —Maximum number of tries to establish the primary path. Range: 0 through 10,000 Default: 0 (The ingress node never stops trying to establish the primary path.)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Connection Between Ingress and Egress Routers on page 229

retry-timer

Syntax	<code>retry-timer <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Amount of time the ingress router waits between attempts to establish the primary path.
Options	<i>seconds</i> —Amount of time between attempts to connect to the primary path. Range: 1 through 600 seconds Default: 30 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Connection Between Ingress and Egress Routers on page 229

revert-timer

Syntax	<code>revert-timer <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. BFD behavior modified in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the amount of time (in seconds) that an LSP must wait before traffic reverts to a primary path. If during this time the primary path experiences any connectivity problem or stability problem, the timer is restarted. If you have configured BFD on the LSP, the Junos OS waits until the BFD session is restored before starting the revert timer counter. If you have configured a value of 0 seconds for the revert-timer statement and traffic is switched to the secondary path, the traffic remains on that path indefinitely. It is never switched back to the primary path unless you intervene.
Options	seconds —Time in seconds. Range: 0 through 65,535 seconds Default: 60 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Revert Timer for LSPs on page 215

rpf-check-policy (Routing Options)

Syntax	<code>rpf-check-policy <i>policy</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Enable you to control whether a reverse path forwarding (RPF) check is performed for a source and group entry before installing a route in the multicast forwarding cache. This makes it possible to use point-to-multipoint LSPs to distribute multicast traffic to Protocol Independent Multicast (PIM) islands situated downstream from the egress routers of the point-to-multipoint LSPs.
Options	<i>policy</i> —Name of the RPF check routing policy.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs on page 306

rsvp-error-hold-time

Syntax	<code>rsvp-error-hold-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	<p>Amount of time MPLS retains RSVP PathErr messages and considers them for CSPF computations. The more time you configure, the more time a source node (ingress of an RSVP LSP) can have to learn about the failures of its LSP by monitoring PathErr messages transmitted from downstream nodes.</p> <p>Information from the PathErr messages is incorporated into subsequent LSP computations, which can improve the accuracy and speed of LSP setup. Some PathErr messages are also used to update traffic engineering database bandwidth information, reducing inconsistencies between the database and the network.</p>
Options	<p>seconds—Amount of time MPLS retains RSVP PathErr messages and considers them for CSPF computations.</p> <p>Range: 0 through 240 seconds</p> <p>Default: 25 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages on page 68

sampling (Protocols MPLS)

Syntax	sampling { cut-off-threshold <i>percentile</i> ; use-average-aggregate; use-percentile <i>percentile</i> ; }
Hierarchy Level	[edit protocols mpls container-label-switched-path <i>lsp-name</i> splitting-merging]
Release Information	Statement introduced in Junos OS Release 14.2.
Description	Configure traffic sampling.
Options	<p>cut-off-threshold <i>percentile</i>—Specify the percentile value to be used as a cut-off threshold in removing outlier bandwidth samples. All the aggregate bandwidth samples determined as outliers are used for computing aggregate bandwidth used at the time of normalization.</p> <p>Default: 0 percentile (the ingress considers all aggregate bandwidth samples for normalization.)</p> <p>Range: 0 through 100</p> <p>use-average-aggregate—Specify the ingress router to take average of the aggregate samples for normalization.</p> <p>This option is mutually exclusive with the use-percentile configuration option.</p> <p>use-percentile <i>percentile</i>—Specify the ingress router to compute and use the pth percentile from all the bandwidth samples, and use that for normalization.</p> <p>This option is mutually exclusive with the use-average-aggregate configuration option.</p> <p>Range: 0 through 100</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• splitting-merging on page 941

secondary (Protocols MPLS)

Syntax	<pre> secondary <i>path-name</i> { adaptive; admin-group { exclude [<i>group-names</i>]; include-all [<i>group-names</i>]; include-any [<i>group-names</i>]; } bandwidth <i>bps</i>; class-of-service <i>cos-value</i>; hop-limit <i>number</i>; no-cspf; no-decrement-ttl; optimize-timer <i>seconds</i>; preference <i>preference</i>; priority <i>setup-priority reservation-priority</i>; (record no-record); retry-limit <i>number</i>; retry-timer <i>seconds</i>; select (manual unconditional); standby; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	<p>Specify one or more secondary paths to use for the LSP. You can configure more than one secondary path. All secondary paths are equal, and the first one that is available is chosen.</p> <p>You can specify secondary paths even if you have not specified any primary paths.</p> <p>Optionally, you can specify preference, CoS, and bandwidth values for the secondary path, which override any equivalent values that you configure for the LSP (at the [edit mpls label-switched-path] hierarchy level).</p>
Options	<p><i>path-name</i>—Name of a path that you created with the path statement.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Primary and Secondary LSPs on page 214

select

Syntax	<code>select (manual unconditional);</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the conditions under which the path is selected to carry traffic. The manual and unconditional options are mutually exclusive.
Options	manual —The path is selected for carrying traffic if it is up and stable for at least the revert timer window (potentially before the revert timer has elapsed). Traffic is sent to other working paths if the current path is down or degraded (receiving errors). unconditional —The path is always selected for carrying traffic, even if it is currently down or degraded (receiving errors).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Specifying the Conditions for Path Selection on page 216

signal-bandwidth

Syntax	<code>signal-bandwidth type;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes], [edit protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the bandwidth encoding of the signal used for path computation and admission control.
Options	type —Configure the type of bandwidth encoding used on the LSP. It can be any of the following values: 10gigether , ds1 , ds3 , e1 , e3 , ethernet , fastether , gigether , stm-1 , stm-4 , stm-16 , stm-64 , stm-256 , sts-1 , vt1-5 , or vt2 .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Signal Bandwidth Type on page 689

smart-optimize-timer

Syntax	<code>smart-optimize-timer seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	<p>Enable the smart optimization timer. When you enable the smart optimization timer on a router, the Junos OS operates on the assumption that the original LSP path is preferable to any alternate or secondary path. When you enable the smart optimization timer and an LSP fails and its traffic is switched to an alternate path, the smart optimization timer starts and waits 3 minutes (this time is configurable). After 3 minutes have passed, the LSP is switched back to the original path. If the original path fails again and the LSP is switched to an alternate path again, the router waits 1 hour before attempting to switch the LSP back to its original path.</p> <p>If you want to disable the smart optimizer, you can set it to zero. The smart-optimize-timer value in seconds indicates the time before which the LSP is switched back to its primary path in case the primary path becomes available. Otherwise, the time to wait is controlled by the optimize-timer, which is usually set to a high value. Some ISPs have the optimize-timer set to once a day. Sometimes after the smart optimizer causes the LSP to be placed back on its primary path, the primary path goes down again within 60 minutes. When this happens, the smart-optimize-timer is disabled automatically, and the optimize-timer (regular path optimization) goes into effect. This is to protect against a flapping link being used.</p>
Default	The smart optimization timer is enabled by default.
Options	<p>seconds—(Optional) Specify the number of seconds to wait before switching an LSP back to its original path. If you do not specify the number of seconds, the default value is used.</p> <p>Range: 0 through 65,535 seconds</p> <p>Default: 180 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Smart Optimize Timer for LSPs on page 255 • Optimizing Signaled LSPs on page 251 • optimize-aggressive on page 913 • optimize-timer on page 916

soft-preemption (Protocols MPLS)

Syntax	soft-preemption;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Attempt to establish a new path for a preempted LSP before tearing it down.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MPLS Soft Preemption on page 238

splitting-merging

Syntax	<pre> splitting-merging { maximum-member-lsps <i>number</i>; maximum-signaling-bandwidth <i>bps</i>; merging-bandwidth <i>bps</i>; minimum-member-lsps <i>number</i>; minimum-signaling-bandwidth <i>bps</i>; normalization; sampling; splitting-bandwidth <i>bps</i>; splitting-merging-threshold <i>percent</i>; }</pre>
Hierarchy Level	[edit protocols mpls container-label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in Junos OS Release 14.2.
Description	Perform splitting and merging.
Options	<p>maximum-member-lsps <i>number</i>—Number of label-switched paths (LSPs) that a container LSP can have as member LSPs at maximum. Default: 1</p> <p>maximum-signaling-bandwidth <i>bandwidth</i>—Amount of bandwidth in bits per second (bps) that can be signaled for an LSP at maximum after normalization. When maximum-signaling-bandwidth is not configured, the value is derived from the splitting-bandwidth. When auto-bandwidth adjustment is done between two normalization events, per LSP auto-bandwidth configuration and thresholds are used instead of the splitting-bandwidth. Default: 1 bps</p> <p>merging-bandwidth <i>bandwidth</i>—Amount of bandwidth in bits per second (bps) that is used for merging during normalization. Default: 1 bps</p> <p>minimum-member-lsps <i>number</i>—Number of LSPd that a container LSP can have as member LSPs at minimum. Default: 64</p> <p>minimum-signaling-bandwidth <i>bandwidth</i>—Amount of bandwidth in bits per second (bps) that can be signaled for an LSP at minimum after normalization. When minimum-signaling-bandwidth is not configured, the value is derived from the merging-bandwidth. When auto-bandwidth adjustment is done between two normalization events, per LSP auto-bandwidth configuration and thresholds are used instead of the merging-bandwidth. Default: 1 bps</p>

splitting-bandwidth *bandwidth*—Amount of bandwidth in bits per second (bps) that can be used for splitting during normalization.

Default: 1 bps

splitting-merging-threshold *percent*—Percentage changes in aggregate bandwidth relevant for splitting and merging.

Default: 0%

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [container-label-switched-path on page 848](#)

srlg

Syntax

```
srlg {  
    srlg-name {  
        srlg-cost srlg-cost;  
        srlg-value srlg-value;  
    }  
}
```

Hierarchy Level [edit routing-options],
[edit logical-systems *logical-system-name* routing-options]
[edit protocols mpls interface *interface-name*]

Release Information Statement introduced in Junos OS Release 11.4.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Configuring Shared Risk Link Group (SRLG) parameters.

Options **srlg-cost *srlg-cost***—Specify a cost for the SRLG ranging from 1 through 65535.
srlg-value *srlg-value*—Specify a Group ID for the SRLG ranging from 1 through 4294967295.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring SRLG on page 81](#)

srlg-cost

Syntax	<code>srlg-cost <i>srlg-cost</i>;</code>
Hierarchy Level	[edit routing-options srlg], [edit logical-systems <i>logical-system-name</i> routing-options srlg]
Release Information	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Specify a cost for the Shared Risk Link Group (SRLG) ranging from 1 through 65535 .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring SRLG on page 81

srlg-value

Syntax	<code>srlg-value <i>srlg-value</i>;</code>
Hierarchy Level	[edit routing-options srlg], [edit logical-systems <i>logical-system-name</i> routing-options srlg]
Release Information	Statement introduced in Junos OS Release 11.4. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Specify a Group ID for the Shared Risk Link Group (SRLG) ranging from 1 through 4294967295 .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring SRLG on page 81

standby

Syntax	standby;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.5 for EX Series switches.
Description	Enable the path to remain up at all times to provide instant switchover if connectivity problems occur.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Hot Standby of Secondary Paths for LSPs on page 267• Configuring Path Protection in an MPLS Network (CLI Procedure)

static-label-switched-path

```
Syntax  static-label-switched-path lsp-name {
        bypass bypass-name {
            bandwidth bps;
            description string;
            next-hop (address | interface-name | address/interface-name);
            push out-label;
            to address;
        }
        ingress {
            bandwidth bps;
            class-of-service cos-value;
            description string;
            install {
                destination-prefix <active>;
            }
            link-protection bypass-name name;
            metric metric;
            next-hop (address | interface-name | address/interface-name);
            node-protection bypass-name name next-next-label label;
            no-install-to-address;
            policing {
                filter filter-name;
                no-auto-policing;
            }
            preference preference;
            push out-label;
            to address;
        }
        transit incoming-label {
            bandwidth bps;
            description string;
            link-protection bypass-name name;
            next-hop (address | interface-name | address/interface-name);
            node-protection bypass-name name next-next-label label;
            pop;
            swap out-label;
        }
    }
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Configure a static LSP.

Options *lsp-name*—Name of the path.

The remaining statements are explained separately.

Required Privilege routing—To view this statement in the configuration.
Level routing-control—To add this statement to the configuration.

Related • [Configuring Static LSPs on page 271](#)
Documentation

statistics (Protocols MPLS)

Syntax	<pre>statistics { auto-bandwidth; file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; interval <i>seconds</i>; no-transit-statistics; traffic-class-statistics; transit-statistics-polling; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p> <p>traffic-class-statistics option introduced in Junos OS Release 14.2.</p>
Description	Enable MPLS statistics collection and reporting.
Options	<p>file <i>filename</i>—(Optional) Name of the file to receive the output. We recommend that you place MPLS tracing output in the file <code>mpls-stat</code> in the <code>/var/log</code> directory.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>file</i> reaches its maximum size, it is renamed <i>file.0</i>, then <i>file.1</i>, and so on, until the maximum number of files is reached. Then, the oldest file is overwritten.</p> <p>Range: 2 or more</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>interval <i>seconds</i>—Interval at which to periodically collect statistics.</p> <p>Range: 1 through 65,535</p> <p>Default: 300 seconds</p> <p>no-world-readable—(Optional) Prevent users from reading the log file.</p> <p>size <i>size</i>—(Optional) Maximum size of each file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a file named <i>file</i> reaches this size, it is renamed <i>file.0</i>. When the <i>file</i> again reaches its maximum size, <i>file.0</i> is renamed <i>file.1</i> and <i>file</i> is renamed <i>file.0</i>. This renaming scheme continues until the maximum number of files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum file size, you also must specify a maximum number of files with the files option.</p> <p>world-readable—(Optional) Enable users to read the log file.</p> <p>Syntax: Syntax: xk to specify KB, xm to specify MB, or xg to specify GB</p> <p>Range: 10 KB through the maximum file size supported on your system</p>

Default: 1 MB

traffic-class-statistics—(Optional) Create counters that maintain data traffic statistics per traffic class at the ingress of all types of LSPs and egress of ultimate hop popping (UHP) point-to-point LSPs. These counters are not created by default and are required to be configured to perform traffic-class-scoped loss measurement.

transit-statistics-polling—(Optional) Enable the polling and display of MPLS statistics for LSPs transiting the router. By default, RSVP does not periodically poll for transit LSP statistics. You cannot configure this statement and the **no-transit-statistics** statement at the same time.

The remaining statements are explained separately.

Required Privilege Level routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- [Configuring MPLS to Gather Statistics on page 342](#)
- [Configuring Automatic Bandwidth Allocation for LSPs on page 257](#)

swap

Syntax	<code>swap out-label;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>], [edit protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Remove the label at the top of the label stack and replace it with the specified label. Manually assigned incoming labels can have values from 1,000,000 through 1,048,575. This statement is used to configure static LSPs at transit routers.
Options	out-label —Manually assigned outgoing label value. Range: 0 through 1,048,575 Default: If you do not define the out-label option, the original label value remains unchanged.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• pop on page 924• push on page 929• Configuring the Intermediate (Transit) and Egress Routers for Static LSPs on page 274

switch-away-lsps

Syntax	switch-away-lsps;
Hierarchy Level	[edit logical-systems <i>logical-systems-name</i> protocols mpls interface <i>interface-name</i>], [edit protocols mpls interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	<p>(MX Series routers only) Enable you to switch an LSP away from a network node using a bypass LSP. This feature could be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic. Configure this statement only after you have configured and committed the always-mark-connection-protection-tlv statement.</p> <p>The always-mark-connection-protection-tlv statement marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. When you configure the switch-away-lsps statement, traffic is switched to the bypass LSP.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Switching LSPs Away from a Network Node on page 479

switching-type

Syntax	switching-type (fiber lambda psc-1 tdm);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes], [edit protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the switching method for the LSP. The switching method can be one of the following values: <ul style="list-style-type: none">• fiber—Fiber switching• lambda—Lambda switching• psc-1—Packet switching• tdm—Time-division multiplexing (TDM) switching
Default	psc-1
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MPLS LSPs for GMPLS on page 688

sync-active-path-bandwidth

Syntax	sync-active-path-bandwidth;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> auto-bandwidth], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>]</p>
Release Information	Statement introduced in Junos OS Release 13.2.
Description	<p>When you have a primary and a secondary path configuration, specify that a path needs to be signaled with the active-path bandwidth when the auto-bandwidth adjustment happens and that the secondary path synchronizes the bandwidth reservations to that of the primary path.</p> <p>When a primary path fails, bandwidth reservations are made by the secondary path on the links that it uses. If you include the sync-active-path-bandwidth statement, the secondary path releases the bandwidth it has reserved and adjusts its bandwidth after the primary path begins carrying traffic.</p> <p>For example, suppose the active path is a secondary path with a reserved bandwidth of 10 GB as a result of the automatic bandwidth adjustment. Then suppose there is a switchover from the secondary path to the primary path. After some time the primary path reserves 5 GB as a result of a new automatic adjustment. Without the sync-active-path-bandwidth statement, the secondary path does not release the 10 GB after a switchover occurs. That bandwidth is wasted. If the sync-active-path-bandwidth is included in the configuration, the secondary path adjusts its bandwidth to 5 GB along with the primary path.</p>
Default	When you have a primary and a secondary path configuration, and the primary path fails, bandwidth reservations are made by the secondary path on the links that it uses. When the primary path comes back and the traffic switches over, the secondary path does not release its bandwidth reservations.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Disabling Constrained-Path LSP Computation on page 239 • Configuring Explicit-Path LSPs on page 278

te-class-matrix

Syntax	<pre>te-class-matrix { tenumber { priority <i>priority</i>; traffic-class { ctnumber priority <i>priority</i>; } } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls diffserv-te], [edit protocols mpls diffserv-te]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the traffic engineering class matrix for a multiclass LSP or a DiffServ-aware traffic engineering LSP.
Default	<p>The default traffic engineering class matrix is:</p> <pre>te-class-matrix { te0 traffic-class ct0 priority 7; te1 traffic-class ct1 priority 7; te2 traffic-class ct2 priority 7; te3 traffic-class ct3 priority 7; te4 traffic-class ct0 priority 0; te5 traffic-class ct1 priority 0; te6 traffic-class ct2 priority 0; te7 traffic-class ct3 priority 0; }</pre> <p>If you define any of the traffic engineering classes, all the default values are dropped.</p>
Options	<p>ctnumber—Specify the number of the class type. It can be one of four values: ct0, ct1, ct2, or ct3.</p> <p>priority <i>priority</i>—Specify the priority of the class type. It can be one of eight values from 0 through 7.</p> <p>tenumber—Specify the number of the traffic engineering class. It can be one of eight values: te0, te1, te2, te3, te4, te5, te6, or te7. You must configure the traffic engineering classes in order, starting with te0.</p> <p>traffic-class—Specify the traffic class for the traffic engineering class.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Traffic Engineering Classes on page 317

to

Syntax	<code>to address;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> bypass],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress],</p> <p>[edit protocols mpls label-switched-path <i>lsp-name</i>],</p> <p>[edit protocols mpls static-label-switched-path <i>lsp-name</i> bypass],</p> <p>[edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	Specify the egress router of a dynamic LSP.
Options	<i>address</i> —Address of the egress router.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Egress Router Address for LSPs on page 212

traceoptions (Protocols MPLS)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series. ted-export option introduced in Junos OS Release 14.2. ted-import option introduced in Junos OS Release 14.2. lsp-history option added in Junos OS Release 15.1.
Description	Configure MPLS tracing options at the protocol level or for a label-switched path. To specify more than one tracing operation, include multiple flag statements.
Default	The default MPLS protocol-level tracing options are inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>filename—Name of the file to receive the output of the tracing operation. All files are placed in the directory /var/log. We recommend that you place MPLS tracing output in the file mpls-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.</p> <p>MPLS Tracing Flags</p> <ul style="list-style-type: none"> • all—Trace all operations • autobw-state—Automatic bandwidth events. • connection—All circuit cross-connect (CCC) activity • connection-detail—Detailed CCC activity

- **cspf**—CSPF computations
- **cspf-link**—Links visited during CSPF computations
- **cspf-node**—Nodes visited during CSPF computations
- **error**—MPLS error packets
- **graceful-restart**—Trace MPLS graceful restart events
- **lsp-history**—Trace LSP history events
- **lsping**—Trace lsping packets and return codes
- **nsr-synchronization**—Trace NSR synchronization events
- **nsr-synchronization-detail**—Trace NSR synchronization events in detail
- **state**—All LSP state transitions
- **static**—Trace static label-switched path
- **ted-export**—Trace leaking of entries from **lsdist.0** table into the traffic engineering database
- **ted-import**—Trace leaking traffic engineering database entries into the **lsdist.0** table
- **timer**—Timer usage

no-world-readable—(Optional) Allow only certain users to read the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Tracing MPLS and LSP Packets and Operations on page 361

traffic-class (delay)

Syntax	<pre>traffic-class <i>tc-value</i> { average-sample-size <i>sample size</i>; padding-size <i>size</i>; query-interval <i>milliseconds</i>; rtt-delay-threshold <i>rtt threshold value</i>; twcd-delay-threshold <i>twcd threshold value</i>; }</pre>
Hierarchy Level	<p>[edit protocols mpls oam performance-monitoring querier delay], [edit protocols mpls label-switched-path <i>lsp-name</i> oam performance-monitoring querier delay], [edit protocols mpls label-switched-path <i>lsp-name</i> primary path-name oam performance-monitoring querier delay], [edit protocols mpls label-switched-path <i>lsp-name</i> secondary path-name oam performance-monitoring querier delay]</p>
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>Configure traffic class specific options.</p> <p>Specify the traffic classes for which loss measurement has to be performed. This parameter takes one of the <i>tc-all tc-0 tc-1 tc-2 tc-3 tc-4 tc-5 tc-6 tc-7 tc-none</i> traffic-class values. For each traffic class, you can configure the respective parameters.</p> <p>To enable traffic-class parameters, configure the traffic-class-statistics configuration statement under the [edit protocol mpls statistic] hierarchy level.</p>
Options	<p>average-sample-size <i>sample size</i>—(Optional) Specify the number of samples used for calculating the average of various metrics. Default: 5 Range: 1 through 30</p> <p>padding-size <i>size</i>—(Optional) Specify the delay-measurement message length, which is used to calculate the delay experienced by messages of different sizes. Default: 0 Range: 1 through 1500</p> <p>query-interval <i>milliseconds</i>—Specify the minimum transmit interval, which signifies how often the loss measurement message is generated from the querier. Default: 10 seconds Range: 1000 through 4294967295 milliseconds</p> <p>rtt-delay-threshold <i>rtt threshold value</i>—Specify the round-trip delay threshold value. Range: 1 through 4294967295 microseconds</p> <p>twcd-delay-threshold <i>twcd threshold value</i>—Specify the two-way channel delay threshold value. Range: 1 through 4294967295 microseconds</p>

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Pro-Active Loss and Delay Measurements on page 448• On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview on page 50• performance-monitoring (Protocols MPLS) on page 922

traffic-class (loss)

Syntax	<pre>traffic-class <i>tc-value</i> { average-sample-size <i>sample size</i>; loss-threshold <i>loss threshold value</i>; loss-threshold-window <i>number of samples for loss threshold</i>; measurement-quantity <i>bytes packets</i>; query-interval <i>milliseconds</i>; }</pre>
Hierarchy Level	<pre>[edit protocols mpls oam performance-monitoring querier loss], [edit protocols mpls label-switched-path <i>lsp-name</i> oam performance-monitoring querier loss], [edit protocols mpls label-switched-path <i>lsp-name</i> primary path-name oam performance-monitoring querier loss], [edit protocols mpls label-switched-path <i>lsp-name</i> secondary path-name oam performance-monitoring querier loss]</pre>
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>Configure traffic class specific options.</p> <p>Specify the traffic classes for which loss measurement has to be performed. This parameter takes one of the <i>tc-all tc-0 tc-1 tc-2 tc-3 tc-4 tc-5 tc-6 tc-7 tc-none</i> traffic-class values. For each traffic class, you can configure the respective parameters.</p> <p>To enable traffic-class parameters, configure the traffic-class-statistics configuration statement under the [edit protocol mpls statistic] hierarchy level.</p>
Options	<p>average-sample-size <i>sample size</i>—(Optional) Specify the number of samples used for calculating the average of various metrics.</p> <p>Default: 5</p> <p>Range: 1 through 30</p> <p>loss-threshold <i>loss threshold value</i>—Specify the threshold value that will be used with loss-threshold-window to calculate the loss within specified window size.</p> <p>Range: 1 through 4294967295</p> <p>loss-threshold-window <i>number of samples for loss threshold</i>—Specify the number of samples used for loss threshold calculation.</p> <p>Range: 1 through 30</p> <p>measurement-quantity <i>bytes packets</i>—(Optional) Specify whether packet or byte loss is being measured at the querier.</p> <p>Default: packets</p> <p>query-interval <i>milliseconds</i>—Specify the minimum transmit interval, which signifies how often the loss measurement message is generated from the querier.</p> <p>Default: 10 seconds</p> <p>Range: 1000 through 4294967295 milliseconds</p>

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Pro-Active Loss and Delay Measurements on page 448• On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview on page 50• performance-monitoring (Protocols MPLS) on page 922

traffic-class (loss-delay)

Syntax	<pre> traffic-class <i>tc-value</i> { average-sample-size <i>sample size</i>; loss-threshold <i>loss threshold value</i>; loss-threshold-window <i>number of samples for loss threshold</i>; measurement-quantity <i>bytes packets</i>; padding-size <i>size</i>; query-interval <i>milliseconds</i>; rtt-delay-threshold <i>rtt threshold value</i>; twcd-delay-threshold <i>twcd threshold value</i>; } </pre>
Hierarchy Level	<p>[edit protocols mpls oam performance-monitoring querier loss-delay], [edit protocols mpls label-switched-path <i>lsp-name</i> oam performance-monitoring querier loss-delay], [edit protocols mpls label-switched-path <i>lsp-name</i> primary <i>path-name</i> oam performance-monitoring querier loss-delay], [edit protocols mpls label-switched-path <i>lsp-name</i> secondary <i>path-name</i> oam performance-monitoring querier loss-delay]</p>
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>Configure traffic class specific options.</p> <p>Specify the traffic classes for which loss measurement has to be performed. This parameter takes one of the <i>tc-all tc-0 tc-1 tc-2 tc-3 tc-4 tc-5 tc-6 tc-7 tc-none</i> traffic-class values. For each traffic class, you can configure the respective parameters.</p> <p>To enable traffic-class parameters, configure the traffic-class-statistics configuration statement under the [edit protocol mpls statistic] hierarchy level.</p>
Options	<p>average-sample-size <i>sample size</i>—(Optional) Specify the number of samples used for calculating the average of various metrics. Default: 5 Range: 1 through 30</p> <p>loss-threshold <i>loss threshold value</i>—Specify the threshold value that will be used with loss-threshold-window to calculate loss within specified window size. Range: 1 through 4294967295</p> <p>loss-threshold-window <i>number of samples for loss threshold</i>—Specify the number of samples used for loss threshold calculation. Range: 1 through 30</p> <p>measurement-quantity <i>bytes packets</i>—(Optional) Specify whether packet or byte loss is being measured at the querier. Default: packets</p> <p>padding-size <i>size</i>—(Optional) Specify the delay-measurement message length, which is used to calculate the delay experienced by messages of different sizes.</p>

Default: 0

Range: 1 through 1500

query-interval *milliseconds*—Specify the minimum transmit interval, which signifies how often the loss measurement message is generated from the querier.

Default: 10 seconds

Range: 1000 through 4294967295 milliseconds

rtt-delay-threshold *rtt threshold value*—Specify the round-trip delay threshold value.

Range: 1 through 4294967295 microseconds

twcd-delay-threshold *twcd threshold value*—Specify the two-way channel delay threshold value.

Range: 1 through 4294967295 microseconds

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Configuring Pro-Active Loss and Delay Measurements on page 448• On-Demand Packet Loss and Delay Measurement for UHP LSPs Overview on page 50• performance-monitoring (Protocols MPLS) on page 922
------------------------------	---

traffic-engineering (Protocols MPLS)

Syntax	traffic-engineering (bgp bgp-igp bgp-igp-both-ribs mpls-forwarding);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Select whether MPLS performs traffic engineering on BGP destinations only or on both BGP and IGP destinations. Affects only LSPs originating from this routing device, not transit or egress LSPs.
Default	bgp
Options	bgp —On BGP destinations only. Ingress routes are installed in the inet.3 routing table. bgp-igp —On both BGP and IGP destinations. Ingress routes are installed in the inet.0 routing table. If IGP shortcuts are enabled, the shortcut routes are automatically installed in the inet.0 routing table. bgp-igp-both-ribs —On both BGP and IGP destinations. Ingress routes are installed in the inet.0 and inet.3 routing tables. This option is used to support VPNs. mpls-forwarding —On both BGP and IGP destinations. Use ingress routes for forwarding only, not for routing.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Traffic Engineering for LSPs on page 336• Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)

traffic-engineering (Protocols BGP)

Syntax	traffic-engineering { unicast; }
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols bgp family],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> family],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family],</p> <p>[edit protocols bgp family],</p> <p>[edit protocols bgp group <i>group-name</i> family],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp family],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family]</p>
Release Information	<p>Statement introduced in Junos OS Release 14.2.</p> <p>Statement introduced for QFX5100 in Junos OS Release 15.1</p>
Description	<p>Enable traffic engineering address family. This generates a multiprotocol address family indicator (AFI) and a subsequent address family identifier (SAFI) to be negotiated with the BGP peers.</p> <p>The BGP network layer reachability information (NLRI) information is exchanged between the peers only when the traffic engineering AFI and SAFI are shared between them. If the peers do not agree on the use of the AFI and SAFI, the connection between the peers is terminated.</p>
Options	unicast —Include BGP-TE NLRI.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Link State Distribution Using BGP on page 364

transit-lsp-association

Syntax	<pre>transit-lsp-association <i>transit-association-lsp-group-name</i> { from-1 <i>address-of-associated-lsp-1</i>; from-2 <i>address-of-associated-lsp-2</i>; lsp-name-1 <i>name-of-associated-lsp-1</i>; lsp-name-2 <i>name-of-associated-lsp-2</i>; }</pre>
Hierarchy Level	[edit protocols mpls]
Release Information	Statement introduced in Junos OS Release 12.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Associate two label-switched paths (LSPs) at a transit node to configure a path for sending and receiving GAL and G-Ach messages for MPLS-TP OAM.
Options	<p><i>transit-association-lsp-group-name</i>—Name of the transit association LSP group.</p> <p><i>from-1 address-of-associated-lsp-1</i>—Address of the first associated LSP.</p> <p><i>from-2 address-of-associated-lsp-2</i>—Address of the second associated LSP.</p> <p><i>lsp-name-1 name-of-associated-lsp-1</i>—Name of the first associated LSP.</p> <p><i>lsp-name-2 name-of-associated-lsp-1</i>—Name of the second associated LSP.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring the MPLS Transport Profile for OAM on page 136

ultimate-hop-popping

Syntax	ultimate-hop-popping;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>label-switched-path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>label-switched-path-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3. Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.
Description	<p>Enable ultimate-hop popping on LSPs. Configure this statement on the device at the LSP ingress. In ultimate-hop popping, the MPLS label is popped from the IP packet at the PE router. The IP address is checked in a second address lookup (also at the PE router), and then the packet is forwarded to its destination.</p> <p>Be aware of the following platform requirements and restrictions:</p> <ul style="list-style-type: none"> • UHP LSPs using VT interfaces—Supported on all M Series, MX Series, T Series, and TX Matrix routers. • UHP LSPs using LSI interfaces—Supported on MX 3D Series routers only. • UHP LSP requirements for the egress PE device—For M Series and T Series routers, a VT interface is needed. • UHP LSPs and Layer 3 VPNs—UHP LSPs are supported for Layer 3 VPNs configured on MX 3D Series routers only. • UHP LSPs and VPLS—UHP LSPs are supported for VPLS configured on MX 3D Series routers only. You must configure the <i>no-tunnel-services</i> statement at the [edit routing-instances <i>routing-instance-name</i> protocols vpls] hierarchy level.
Default	Ultimate-hop popping is disabled by default on LSPs. Penultimate-hop popping is the default behavior. In penultimate-hop popping, the final MPLS label is popped from the IP packet at the last provider router in the network before being forwarded to the PE router. The PE router receives the packet and checks the IP address, and then the packet is forwarded to its destination.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Ultimate-Hop Popping for LSPs on page 222 • explicit-null on page 865

RSVP Configuration Statements

- [\[edit protocols rsvp\] Hierarchy Level on page 968](#)
- [\[edit protocols rsvp\] Hierarchy Level on page 971](#)
- [admin-group on page 973](#)
- [aggregate \(Protocols RSVP\) on page 974](#)
- [authentication-key \(Protocols RSVP\) on page 975](#)
- [bandwidth \(Protocols RSVP\) on page 976](#)
- [bypass \(Signaled LSP\) on page 977](#)
- [bypass \(Static LSP\) on page 978](#)
- [class-of-service \(Protocols RSVP\) on page 979](#)
- [destination-networks on page 980](#)
- [devices on page 981](#)
- [disable \(Protocols RSVP\) on page 982](#)
- [dynamic-bidirectional-transport on page 983](#)
- [fast-reroute \(Protocols RSVP\) on page 983](#)
- [graceful-deletion-timeout on page 984](#)
- [graceful-restart \(Protocols RSVP\) on page 985](#)
- [hello-acknowledgements on page 986](#)
- [hello-interval \(Protocols RSVP\) on page 986](#)
- [hop-limit on page 987](#)
- [interface \(Protocols RSVP\) on page 988](#)
- [keep-multiplier on page 989](#)
- [label-switched-path-template \(Multicast\) on page 990](#)
- [link-protection \(RSVP\) on page 992](#)
- [load-balance \(Protocols RSVP\) on page 993](#)
- [max-bypasses on page 994](#)
- [no-local-reversion on page 995](#)
- [node-hello on page 996](#)
- [no-adjacency-down-notification \(Protocols IS-IS\) on page 997](#)

- [no-cspf \(Protocols RSVP\) on page 998](#)
- [no-interface-hello on page 998](#)
- [no-neighbor-down-notification on page 999](#)
- [no-node-id-subobject on page 999](#)
- [no-p2mp-sublsp on page 1000](#)
- [node-link-protection \(Protocols MPLS\) on page 1000](#)
- [optimize-timer \(Protocols RSVP\) on page 1001](#)
- [path \(Protocols RSVP\) on page 1002](#)
- [peer-interface \(Protocols RSVP\) on page 1003](#)
- [preemption on page 1004](#)
- [priority \(Protocols RSVP\) on page 1005](#)
- [refresh-time on page 1006](#)
- [reliable on page 1006](#)
- [rsvp on page 1007](#)
- [rsvp-te \(Routing Options\) on page 1008](#)
- [setup-protection on page 1008](#)
- [soft-preemption \(Protocols RSVP\) on page 1009](#)
- [static-label-switched-path on page 1010](#)
- [subscription on page 1011](#)
- [traceoptions \(Protocols RSVP\) on page 1012](#)
- [transit on page 1014](#)
- [tunnel-services \(RSVP\) on page 1015](#)
- [ultimate-hop-popping on page 1016](#)
- [update-threshold on page 1017](#)

[\[edit protocols rsvp\] Hierarchy Level](#)

The following statement hierarchy can also be included at the [\[edit logical-systems *logical-system-name*\]](#) hierarchy level.

```
protocols {  
  rsvp {  
    disable;  
    fast-reroute optimize-timer seconds;  
    graceful-deletion-timeout seconds;  
    graceful-restart {  
      disable;  
      helper-disable;  
      maximum-helper-recovery-time seconds;  
      maximum-helper-restart-time seconds;  
    }  
    interface interface-name {  
      ... the interface subhierarchy appears after the main [edit protocols rsvp] hierarchy ...  
    }  
  }  
}
```

```

keep-multiplier number;
load-balance bandwidth;
node-hello;
no-interface-hello;
no-node-id-subobject;
peer-interface peer-interface-name {
    disable;
    (aggregate | no-aggregate);
    authentication-key key;
    hello-interval seconds;
    (reliable | no-reliable);
}
preemption {
    (aggressive | disabled | normal);
    soft-preemption cleanup-timer seconds;
}
refresh-time seconds;
traceoptions {
    file filename <files number> <size maximum-file-size> <world-readable |
    no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
tunnel-services {
    devices device-names;
}
}

rsvp {
    interface interface-name {
        disable;
        (aggregate | no-aggregate);
        authentication-key key;
        bandwidth bps;
        hello-interval seconds;
        link-protection {
            ... the link-protection subhierarchy appears after the main [edit protocols rsvp
            interface interface-name] hierarchy ...
        }
        (reliable | no-reliable);
        subscription {
            percentage;
            ct0 percentage;
            ct1 percentage;
            ct2 percentage;
            ct3 percentage;
        }
        update-threshold percentage;
    }
}

interface interface-name {
    link-protection {
        disable;
        admin-group {
            exclude [ group-names ];
            include-all [ group-names ];
            include-any [ group-names ];
        }
    }
}

```

```

    }
    bandwidth {
        bps;
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
    bypass bypass-name {
        ... the bypass subhierarchy appears after the main [edit protocols rsvp interface
        interface-name link-protection] hierarchy ...
    }
    class-of-service cos-value;
    hop-limit number;
    max-bypasses number;
    no-cspf;
    no-node-protection;
    optimize-timer seconds;
    path address <loose| strict>;
    priority setup-priority reservation-priority;
    subscription percentage;
}

link-protection {
    bypass bypass-name {
        admin-group {
            exclude [ group-names ];
            include-all [ group-names ];
            include-any [ group-names ];
        }
        bandwidth {
            bps;
            ct0 bps;
            ct1 bps;
            ct2 bps;
            ct3 bps;
        }
        class-of-service cos-value;
        description text;
        hop-limit number;
        no-cspf;
        path address <loose| strict>;
        priority setup-priority reservation-priority;
        to address;
    }
}
}
}
}
}

```

Related Documentation

- *Notational Conventions Used in Junos OS Configuration Hierarchies*
- *[edit protocols] Hierarchy Level*

[edit protocols rsvp] Hierarchy Level

The following statements can also be configured at the **[edit logical-systems *logical-system-name*]** hierarchy level:

```

protocols {
  rsvp {
    disable;
    fast-reroute optimize-timer seconds;
    graceful-deletion-timeout seconds;
    graceful-restart {
      disable;
      helper-disable;
      maximum-helper-recovery-time seconds;
      maximum-helper-restart-time seconds;
    }
    interface interface-name {
      disable;
      (aggregate | no-aggregate);
      authentication-key key;
      bandwidth bps;
      hello-interval seconds;
      link-protection {
        disable;
        admin-group {
          exclude group-names;
          include-all group-names;
          include-any group-names;
        }
        bandwidth bandwidth;
        bypass bypass-name {
          bandwidth bps {
            ct0 bps;
            ct1 bps;
            ct2 bps;
            ct3 bps;
          }
          description text;
          hop-limit number;
          no-cspf;
          path address <strict | loose>;
          priority setup-priority reservation-priority;
          to address;
        }
        class-of-service cos-value;
        exclude-srlg;
        hop-limit number;
        max-bypasses number;
        no-cspf;
        no-node-protection;
        optimize-timer seconds;
        path address <strict | loose>;
        priority setup-priority reservation-priority;
        subscription percentage {

```

```

        ct0 percentage;
        ct1 percentage;
        ct2 percentage;
        ct3 percentage;
    }
}
(reliable | no-reliable);
subscription percentage {
    ct0 percentage;
    ct1 percentage;
    ct2 percentage;
    ct3 percentage;
}
update-threshold percentage;
}
keep-multiplier number;
load-balance {
    bandwidth;
}
no-node-id-subobject;
no-p2mp-sublsp;
peer-interface peer-interface-name {
    (aggregate | no-aggregate);
    authentication-key key;
    disable;
    hello-interval seconds;
    (reliable | no-reliable);
}
preemption {
    (aggressive | disabled | normal);
    soft-preemption {
        cleanup-timer seconds;
    }
}
refresh-time seconds;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
tunnel-services {
    devices device-names;
}
}
}

```

admin-group

Syntax	<pre>admin-group { exclude [<i>group-names</i>]; include-all [<i>group-names</i>]; include-any [<i>group-names</i>]; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</pre>
Release Information	Statement introduced in Junos OS Release 9.2.
Description	<p>Enable you to configure administrative groups for bypass label-switched paths (LSPs). You can configure administrative groups either globally for all bypass LSPs traversing an interface or for just a specific bypass LSP.</p>
Options	<p>exclude <i>group-names</i>—Specify the administrative groups to exclude for a bypass LSP.</p> <p>include-all <i>group-names</i>—Specify the administrative groups whose links the bypass LSP must traverse.</p> <p>include-any <i>group-names</i>—Specify the administrative groups whose links the bypass LSP can traverse.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Administrative Groups for Bypass LSPs on page 504

aggregate (Protocols RSVP)

Syntax	(aggregate no-aggregate);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	<p>Control the use of RSVP aggregate messages on an interface or peer interface:</p> <ul style="list-style-type: none">• aggregate—Use RSVP aggregate messages.• no-aggregate—Do not use RSVP aggregate messages. <p>Aggregate messages can pack multiple RSVP messages into a single transmission, thereby reducing network overhead and enhancing efficiency. The number of supportable sessions and processing overhead are significantly improved when aggregation is enabled.</p> <p>Not all routers connected to a subnet need to support aggregation simultaneously. Each RSVP router negotiates its intention to use aggregate messages on a per-neighbor basis. Only when both routers agree are aggregate messages sent.</p> <p>To have refresh reduction and reliable delivery, you must include the aggregate and reliable statements.</p>
Default	Aggregation is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RSVP Refresh Reduction on page 473• reliable on page 1006

authentication-key (Protocols RSVP)

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	<p>Authentication key (password). Neighboring routers use the password to verify the authenticity of packets sent from this interface or peer interface.</p> <p>RSVP uses HMAC-MD5 authentication, which is defined in RFC 2104, <i>HMAC: Keyed-Hashing for Message Authentication</i>.</p> <p>All routers that are connected to the same IP subnet must use the same authentication scheme and password.</p>
Options	key —Authentication password. It can be 1 through 16 contiguous digits or letters. Separate decimal digits with periods. Separate hexadecimal digits with periods and precede the string with 0x. If you include spaces in the password, enclose the entire password in quotation marks (" ").
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring RSVP Authentication on page 476

bandwidth (Protocols RSVP)

Syntax	<code>bandwidth <i>bps</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i></code> <code>link-protection],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i></code> <code>link-protection bypass <i>bypass-name</i>],</code> <code>[edit protocols rsvp interface <i>interface-name</i>],</code> <code>[edit protocols rsvp interface <i>interface-name</i> link-protection],</code> <code>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	<p>For certain logical interfaces (such as Asynchronous Transfer Mode [ATM], Permanent Virtual Circuit [PVC], or Frame Relay), you cannot determine the correct bandwidth from the hardware. This statement enables you to specify the actual available bandwidth.</p> <p>This statement also enables you to specify the bandwidth for a bypass label switched path (LSP). If you have configured multiple bypasses, this statement is mandatory and is applied to all of the bypass LSPs.</p>
Default	The hardware raw bandwidth is used.
Options	<p><i>bps</i>—Bandwidth in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations k (for a thousand), m (for a million), or g (for a billion [also called a thousand million]).</p> <p>Range: Any positive integer</p> <p>Default: 0 (no bandwidth is reserved)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring the Bandwidth for Bypass LSPs on page 504• Configuring Link Protection on Interfaces Used by LSPs on page 502• Configuring Bypass LSPs on page 503

bypass (Signaled LSP)

Syntax	<pre>bypass <i>bypass-name</i> { bandwidth <i>bps</i>; description <i>text</i>; hop-limit <i>number</i>; no-cspf; path <i>address</i> <strict loose>; priority <i>setup-priority reservation-priority</i>; to <i>address</i>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>The description option was added in Junos OS Release 10.4.</p>
Description	<p>Enables you to configure specific bandwidth and path constraints for a bypass LSP. It is possible to individually configure multiple bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints.</p> <p>If you specify the bandwidth, hop-limit, and path statements for the bypass LSP, these values take precedence over the values configured at the [edit protocols rsvp interface <i>interface-name</i> link-protection] hierarchy level. The other attributes (subscription, no-node-protection, and optimize-timer) are inherited from the general constraints.</p>
Options	<p>bypass-name—(Required) Specify a name for the bypass LSP. The name can be up to 64 characters.</p> <p>description—Provides a textual description of the bypass LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the show mpls lsp bypass detail command and has no effect on the operation of the bypass LSP. The description text can be no more than 80 characters in length.</p> <p>to address—(Required) Specify the address for the interface of the immediate next-hop node (for link protection) or the next-next-hop node (for node-link protection). The address specified determines whether this is a link protection bypass or a node-link protection bypass. On multiaccess networks (for example, a LAN), this address is also used to specify which next-hop node is being protected.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Bypass LSPs on page 503

bypass (Static LSP)

Syntax	<pre>bypass <i>bypass-name</i> { bandwidth <i>bps</i>; description <i>string</i>; next-hop (<i>address</i> <i>interface-name</i> <i>address/interface-name</i>); push <i>out-label</i>; to <i>address</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i>], [edit protocols mpls static-label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 10.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	<p>Configure specific bandwidth and path constraints for a bypass ingress LSP. It is possible to configure multiple bypass LSPs individually. If you do not, they all share the same path and bandwidth constraints.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Static LSPs on page 271

class-of-service (Protocols RSVP)

Syntax	<code>class-of-service <i>cos-value</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Class-of-service (CoS) value given to all packets in the bypass LSP. You can specify a single CoS value for all the bypass LSPs traversing an interface. You can also configure CoS values for specific bypass LSPs traversing an interface.</p> <p>The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.</p>
Options	<p><i>cos-value</i>—CoS value. A higher value typically corresponds to a higher level of service.</p> <p>Range: 0 through 7</p> <p>Default: If you do not specify a CoS value, the IP precedence bits from the packet's IP header are used as the packet's CoS value.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Class of Service for Bypass LSPs on page 505

destination-networks

Syntax	<code>destination-networks <i>prefix</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry</i>], [edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry</i>], [edit routing-options dynamic-tunnels <i>tunnel-name</i>], [edit routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Specify the IPv4 prefix range for the destination network. Only tunnels within the specified IPv4 prefix range can be created.
Options	<i>prefix</i> —Destination prefix of the network.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring GRE Tunnels for Layer 3 VPNs</i>• <i>Configuring Dynamic Tunnels</i>• Configuring RSVP Automatic Mesh on page 482

devices

Syntax	<code>devices <i>device-names</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	Specifies one of the virtual tunnel (VT) interfaces to de-encapsulate the egress traffic for ultimate-hop popping on point-to-multipoint LSPs. If no device is specified, the selection process is performed automatically.
Default	The device selection process is performed automatically if no device is configured. Junos OS selects one of the available VT interfaces to de-encapsulate the egress traffic.
Options	<i>device-names</i> —Specify which VT interfaces are used to handle the RSVP traffic. Range: 0 to 8 devices
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs on page 490 • Understanding Redundant Virtual Tunnel Interfaces in MBGP MVPNs

disable (Protocols RSVP)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit logical-systems <i>logical-system-name</i> protocols rsvp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp], [edit protocols rsvp graceful-restart], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Explicitly disable RSVP or RSVP graceful restart. Explicitly disable link protection on the specified interface.
Default	RSVP is enabled on interfaces and peer interfaces configured with the RSVP interface statement. RSVP graceful restart is enabled on the router. Link protection is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Minimum RSVP Configuration on page 471• Configuring RSVP Graceful Restart on page 514• Configuring Link Protection on Interfaces Used by LSPs on page 502

dynamic-bidirectional-transport

Syntax	<code>dynamic-bidirectional-transport { template <i>template</i>; }</code>
Hierarchy Level	<code>[edit protocols rsvp peer-interface <i>peer-interface-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 14.2.
Description	Enable the dynamic setup of associated bidirectional packet LSP for transporting non-packet Generalized Multiprotocol Label Switching (GMPLS) label-switched path (LSP).
Options	template <i>template</i> —Name of the template for the dynamic bidirectional packet LSP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

fast-reroute (Protocols RSVP)

Syntax	<code>fast-reroute optimize-timer <i>seconds</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]</code>
Release Information	Statement added in Junos OS Release 7.5. Statement introduced in Junos OS Release 14.1 for the QFX Series.
Description	Configure the optimize timer for fast reroute. The optimize timer triggers a periodic optimization process that recomputes the fast reroute detour LSPs to use network resources more efficiently.
Options	<i>seconds</i> —Specify the number of seconds between fast reroute detour LSP optimizations. Range: 0 through 65,535 seconds Default: 0 (disabled)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Optimization Interval for Fast Reroute Paths on page 228

graceful-deletion-timeout

Syntax	<code>graceful-deletion-timeout seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the time, in seconds, before completing graceful deletion of signaling.
Options	seconds —Time before completing graceful deletion of signaling. Range: 1 through 300 seconds Default: 30 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Graceful Deletion Timeout Interval on page 691

graceful-restart (Protocols RSVP)

Syntax	<pre>graceful-restart { disable; helper-disable; maximum-helper-recovery-time <i>seconds</i>; maximum-helper-restart-time <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options], [edit protocols rsvp], [edit routing-options]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure graceful restart on the router. You must configure the graceful-restart statement at the [edit routing-options] hierarchy level to enable graceful restart on the router.
Options	<p>disable—Disable graceful restart on the router or for RSVP.</p> <p>helper-disable—Disable RSVP graceful restart helper mode (this option is only available at the [edit protocols rsvp] hierarchy level).</p> <p>Default: Helper mode is enabled by default.</p> <p>maximum-helper-recovery-time <i>seconds</i>—The maximum length of time the router stores the state of neighboring routers when they undergo a graceful restart. The value applies to all neighboring routers, so it should be based on the time that the slowest RSVP neighbor requires for restart.</p> <p>Default: 180 seconds</p> <p>Range: 1 through 3600 seconds</p> <p>maximum-helper-restart-time <i>seconds</i>—The maximum length of time the router waits between when it discovers that a neighboring router has gone down and when it declares the neighbor down. This value is applied to all neighboring routers, so it should be based on the time that the slowest RSVP neighbor requires for restart.</p> <p>Default: 20 seconds</p> <p>Range: 1 through 1800 seconds</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring RSVP Graceful Restart on page 514

hello-acknowledgements

Syntax	hello-acknowledgements;
Hierarchy Level	[edit logical-systems <i>logical-systems-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in Junos OS Release 10.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable hello messages from nonsession neighbors to be acknowledged with a hello acknowledgment message. Once hello acknowledgments are enabled, the router continues to acknowledge hello messages from any nonsession RSVP neighbors unless the interface itself goes down or the configuration is changed by an administrator.
Default	Hello acknowledgments are disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Hello Acknowledgments for Nonsession RSVP Neighbors on page 479

hello-interval (Protocols RSVP)

Syntax	hello-interval <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable the sending of hello packets on the interface.
Options	<i>seconds</i> —Length of time between hello packets. A value of 0 disables the sending of hello packets on the interface. Range: 1 through 60 seconds Default: 9 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the RSVP Hello Interval on page 475

hop-limit

Syntax	<code>hop-limit <i>number</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i> fast-reroute], [edit protocols mpls label-switched-path <i>lsp-name</i> (primary secondary) <i>path-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	<p>Specify the maximum number of routers that an LSP can traverse. This limit can be applied to any of the following:</p> <ul style="list-style-type: none"> • LSPs—The configured hop limit includes the ingress and egress routers. You can specify a hop limit for an LSP and for both primary and secondary paths. • Fast reroute detour—Specify the number of additional routers a fast reroute detour can traverse relative to the protected LSP. For example, if an LSP traverses 4 routers, any detour for the LSP can be no more than 10 router hops, including the ingress and egress routers. • Link protection bypass—Specify the maximum number of routers that a link protection bypass can traverse.
Options	<p><i>number</i>—Maximum number of hops.</p> <p>Range: 2 through 255 (for an LSP or for a link protection bypass); 0 through 255 (for fast reroute)</p> <p>Default: 255 (for an LSP or for a link protection bypass); 6 (for fast reroute)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Fast Reroute on page 226 • Limiting the Number of Hops in LSPs on page 256 • Configuring the Hop Limit for Bypass LSPs on page 505

interface (Protocols RSVP)

```
Syntax  interface interface-name {
        disable;
        (aggregate | no-aggregate);
        authentication-key key;
        bandwidth bps;
        hello-interval seconds;
        link-protection {
            disable;
            admin-group {
                exclude [ group-names ];
                include-all [ group-names ];
                include-any [ group-names ];
            }
            bandwidth bps;
            bypass bypass-name {
                bandwidth bps {
                    ct0 bps;
                    ct1 bps;
                    ct2 bps;
                    ct3 bps;
                }
                description text;
                class-of-service cos-value;
                hop-limit number;
                no-cspf;
                path address <strict | loose>;
                priority setup-priority reservation-priority;
                to address;
            }
            class-of-service cos-value;
            hop-limit number;
            max-bypasses number;
            no-cspf;
            no-node-protection;
            optimize-timer seconds;
            path address <strict | loose>;
            priority setup-priority reservation-priority;
            subscription percentage;
        }
        (reliable | no-reliable);
        subscription percentage {
            ct0 percentage;
            ct1 percentage;
            ct2 percentage;
            ct3 percentage;
        }
        update-threshold threshold;
    }
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
[edit protocols rsvp]

Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable RSVP on one or more router interfaces.
Default	RSVP is disabled on all interfaces.
Options	<i>interface-name</i> —Name of an interface. To configure all interfaces, specify all . For details about specifying interfaces, see the <i>Junos OS Network Interfaces Library for Routing Devices</i> . The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Minimum RSVP Configuration on page 471

keep-multiplier

Syntax	keep-multiplier <i>number</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Set the keep multiplier value.
Options	<i>number</i> —Multiplier value. Range: 1 through 255 Default: 3
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Timers for RSVP Refresh Messages on page 483

label-switched-path-template (Multicast)

Syntax	<pre>label-switched-path-template { (default-template <i>lsp-template-name</i>); }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel rsvp-te],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel ingress-replication label-switched-path],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i> rsvp-te],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry-name</i>],</p> <p>[edit protocols mvpn inter-region-segmented template <i>template-name</i> region <i>region-name</i> ingress-replication label-switched-path],</p> <p>[edit protocols mvpn inter-region-segmented template <i>template-name</i> region <i>region-name</i> rspe-te],</p> <p>[edit protocols mvpn inter-region-template template <i>template-name</i> all-regions ingress-replication label-switched-path],</p> <p>[edit protocols mvpn inter-region-template template <i>template-name</i> all-regions rsvp-te],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel ingress-replication label-switched-path],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel rsvp-te],</p> <p>[edit routing-instances <i>routing-instance-name</i> provider-tunnel selective group <i>address</i> source <i>source-address</i> rsvp-te],</p> <p>[edit routing-options dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry-name</i>]</p>
Release Information	Statement introduced in Junos OS Release 8.5.
Description	<p>Specify the LSP template. An LSP template is used as the basis for other dynamically generated LSPs. This feature can be used for a number of applications, including point-to-multipoint LSPs, flooding VPLS traffic, configuring ingress replication for IP multicast using MBGP MVPNs, and to enable RSVP automatic mesh. There is no default setting for the label-switched-path-template statement, so you must configure either the default-template using the default-template option, or you must specify the name of your preconfigured LSP template.</p>
Options	<p>default-template—Specify that the default LSP template be used for the dynamically generated LSPs.</p> <p><i>lsp-template-name</i>—Specify the name of an LSP to be used as a template for the dynamically generated LSPs.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs</i> • <i>Configuring Point-to-Multipoint LSPs for an MBGP MVPN</i> • <i>Flooding Unknown Traffic Using Point-to-Multipoint LSPs in VPLS</i>

- [Configuring RSVP Automatic Mesh on page 482](#)

link-protection (RSVP)

Syntax	<pre> link-protection { disable; admin-group { exclude [group-names]; include-all [group-names]; include-any [group-names]; } bandwidth bps; bypass bypass-name { bandwidth bps { ct0 bps; ct1 bps; ct2 bps; ct3 bps; } description text; class-of-service cos-value; hop-limit number; no-cspf; path address <strict loose>; priority setup-priority reservation-priority; to address; } class-of-service cos-value; hop-limit number; max-bypasses number; no-cspf; no-node-protection; optimize-timer seconds; path address <strict loose>; priority setup-priority reservation-priority; subscription percentage; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.
Description	Enable link protection on the specified interface. Using link protection, you can configure a network to reroute traffic quickly around broken links. To fully enable link protection, you also need to configure the link-protection statement at the [edit protocols mpls label-switched-path <i>lsp-name</i>] hierarchy level. You can configure single or multiple bypasses for protected interface.
Default	Link protection is disabled.

Options **no-node-protection**—Disable node-link protection on the RSVP interface. Link protection remains active. When this option is configured, the router can only initiate a next-hop bypass, not a next-next-hop bypass.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation • [Configuring Link Protection on Interfaces Used by LSPs on page 502](#)
 • [link-protection \(Dynamic LSPs\) on page 888](#)

load-balance (Protocols RSVP)

Syntax load-balance {
 bandwidth;
 }

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp],
 [edit protocols rsvp]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Load-balance traffic between RSVP LSPs.

Options **bandwidth**—Load-balance traffic between RSVP LSPs based on the bandwidth configured for each LSP.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation • [Configuring Load Balancing Across RSVP LSPs on page 481](#)

max-bypasses

Syntax	<code>max-bypasses <i>number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Range modified in Junos OS Release 9.3.
Description	Specify the maximum number of dynamic bypass LSPs permitted for protecting this interface. When this option is configured, multiple bypasses for link protection are enabled. Call admission control (CAC) is also enabled. The limit on bypasses configured applies only to dynamically generated bypass LSPs. By default, this option is disabled and only one dynamic bypass LSP is enabled for each interface. If you configure max-bypasses , you must also configure the bandwidth statement.
Options	number —Configure the maximum number of bypass LSPs. If you configure a value of 0, no dynamic bypass LSPs are allowed to be established for the interface. Only static bypass LSPs can be configured. Range: 0 through 99 Default: 1
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Maximum Number of Bypass LSPs on page 506

no-local-reversion

Syntax	no-local-reversion;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	<p>Disables RSVP local revertive mode as specified in RFC 4090, <i>Fast Reroute Extensions to RSVP-TE for LSP Tunnels</i>. RSVP local revertive mode is supported on all Juniper Networks routers running the Junos OS. It is the default behavior. If you include this statement, the Juniper Networks router uses global revertive mode instead. You might need to disable RSVP local revertive mode on Juniper Networks routers if your network includes equipment that does not support this mode.</p> <p>The following information can also be found in RFC 4090. Refer to the full RFC for additional information. When an LSP fails, the connection can be repaired locally using a traffic protection mechanism such as fast reroute. To restore the LSP to a full working path, RFC 4090 specifies the following strategies:</p> <ul style="list-style-type: none"> • Local revertive mode—Upon detecting that the path is restored, the point of local repair (PLR) resignals each of the LSPs that were formerly routed over the restored path. Every LSP successfully resigaled along the restored path is switched back. • Global revertive mode—The ingress router of each tunnel is responsible for reoptimizing the LSPs that used the failed path. There are several potential reoptimization triggers: RSVP error messages, inspection of OSPF LSAs or IS-IS LSPs, and timers. This re-optimization process can proceed as soon as the failure is detected. It is not tied to the restoration of the failed path.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

node-hello

Syntax	node-hello;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in JUNOS Release 10.0. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enables node-ID based RSVP hellos globally on all of the RSVP interfaces on the router to allow Juniper Networks routers to interoperate with the equipment of other vendors. By default, the JUNOS Software uses interface-based RSVP hellos and node-ID based RSVP hellos are disabled. If you have not enabled RSVP node IDs on the router, the JUNOS software does not accept any node-ID hello packets.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RSVP Node ID Hellos on page 478

no-adjacency-down-notification (Protocols IS-IS)

Syntax	no-adjacency-down-notification;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols isis interface <i>interface-name</i>], [edit protocols isis interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.0.
Description	<p>Disable adjacency down notification for IS-IS to allow for migration from IS-IS to OSPF without disruption of the RSVP neighbors and associated RSVP-signaled label-switched paths (LSPs).</p> <p>Whenever IS-IS is deactivated, the IS-IS adjacencies are brought down. IS-IS signals to RSVP to bring down any RSVP neighbors associated with the IS-IS adjacencies, and this further causes the associated LSPs signaled by RSVP to go down as well.</p> <p>A similar process occurs whenever OSPF is deactivated. The OSPF neighbors are brought down. OSPF signals to RSVP to bring down any of the RSVP neighbors associated with the OSPF neighbors, and this further causes the associated LSPs signaled by RSVP to go down as well.</p> <p>If you need to migrate from IS-IS to OSPF or from OSPF to IS-IS, the internal gateway protocol (IGP) notification to RSVP for an adjacency or neighbor down event needs to be ignored. Using the no-adjacency-down-notification or no-neighbor-down-notification statements, you can disable IS-IS adjacency down notification or OSPF neighbor down notification, respectively, until the migration is complete. The network administrator is responsible for configuring the statements before the migration, and then removing them from the configuration afterward, so that IGP notification can function normally.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • no-neighbor-down-notification on page 999

no-cspf (Protocols RSVP)

Syntax	no-cspf;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]
Release Information	Statement introduced in Junos OS Release 7.5.
Description	Disable CSPF computation on all bypass LSPs or on a specific bypass LSP. You need to disable CSPF for link protection to function properly on interarea paths.
Default	CSPF is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Disabling CSPF for Bypass LSPs on page 506

no-interface-hello

Syntax	no-interface-hello;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in JUNOS Release 10.0. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Allows you to explicitly disable RSVP interface hellos globally on the router. This type of configuration might be necessary in networks where the Juniper Networks router has numerous RSVP connections with equipment from other vendors. However, if you disable RSVP interface hellos globally, you can also configure a hello interval on an RSVP interface using the hello-interval (Protocols RSVP) statement. This configuration disables RSVP interface hellos globally but enables RSVP interface hellos on the specified interface. This configuration might be necessary in a heterogeneous network where some devices support RSVP node ID hellos and other devices support RSVP interface hellos.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring RSVP Node ID Hellos on page 478 • hello-interval (Protocols RSVP) on page 986

no-neighbor-down-notification

Syntax	no-neighbor-down-notification;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i>], [edit protocols ospf area <i>area-id</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 8.0.
Description	Disable neighbor down notification for OSPF to allow for migration from OSPF to IS-IS without disruption of the RSVP neighbors and associated RSVP-signaled LSPs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-node-id-subobject

Syntax	no-node-id-subobject;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Disable the record route object (RRO) node ID subobject for compatibility with earlier versions of the Junos OS. To interoperate with other vendors' equipment, the Junos OS supports the RRO node ID subobject for use in inter-AS link and node protection configurations.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Inter-AS Node and Link Protection on page 510

no-p2mp-sublsp

Syntax	no-p2mp-sublsp;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Reject Resv messages that include the S2L_SUB_LSP object. By default, Resv messages that include the S2L_SUB_LSP object are accepted. However, in a network which includes Juniper Networks devices running both Junos OS Release 9.2 and later and Junos OS Release 9.1 and earlier, it is necessary to configure the no-p2mp-sublsp statement on devices running Junos OS Release 9.2 and later to ensure that point-to-multipoint LSPs function properly.
Default	Resv messages that include the S2L_SUB_LSP object are accepted.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases on page 308

node-link-protection (Protocols MPLS)

Syntax	node-link-protection;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i>], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.
Description	Enable node and link protection on the specified LSP. To fully enable node and link protection, you also need to include the link-protection statement at the [edit protocols rsvp interface <i>interface-name</i>] hierarchy level.
Default	Node and link protection is disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Node Protection or Link Protection for LSPs on page 509• MPLS Feature Support on QFX Series and EX4600 Switches• Interprovider and Carrier-of-Carriers VPNs

optimize-timer (Protocols RSVP)

Syntax	<code>optimize-timer <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure an optimize timer for a bypass LSP. The optimize timer initiates a periodic optimization process that reshuffles data LSPs among bypass LSPs to achieve the most efficient use of network resources. The optimization process attempts to either minimize the number of bypasses currently in use, minimize the total amount of bandwidth reserved for all bypasses, or both.
Options	<p><i>seconds</i>—Specify the number of seconds between optimizations.</p> <p>Range: 0 through 65,535 seconds</p> <p>Default: 0 (disabled)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Optimization Interval for Bypass LSPs on page 507

path (Protocols RSVP)

Syntax	<code>path address <strict loose>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure an explicit path (a sequence of strict or loose routes) to control where and how a bypass LSP is established. If multiple bypasses are configured, they all will use the same explicit path.
Default	No path is configured. CSPF automatically calculates the path the bypass LSP takes.
Options	<p>address—IP address of each transit router in the LSP. You must specify the address or hostname of each transit router, although you do not need to list each transit router if its type is loose. As an option, you can include the ingress and egress routers in the path. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path until reaching the egress router (optional) or the router immediately before the egress router.</p> <p>Default: If you do not specify any routers explicitly, no routing limitations are imposed on the bypass LSP.</p> <p>loose—(Optional) The next address in the path statement is loose. The LSP can traverse other routers before reaching this router.</p> <p>Default: strict</p> <p>strict—(Optional) The LSP must go to the next address specified in the path statement without traversing other nodes. This is the default.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring an Explicit Path for Bypass LSPs on page 508

peer-interface (Protocols RSVP)

Syntax	<pre>peer-interface <i>peer-interface-name</i> { disable; (aggregate no-aggregate); authentication-key <i>key</i>; dynamic-bidirectional-transport template <i>template</i>; hello-interval <i>seconds</i>; (reliable no-reliable); }</pre>
Hierarchy Level	[edit protocols rsvp]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>dynamic-bidirectional-transport template <i>template</i> option introduced in Junos OS Release 14.2.</p>
Description	<p>Configure the name of the LMP peer device.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring RSVP and OSPF for LMP Peer Interfaces on page 685

preemption

Syntax	<pre>preemption { (aggressive disabled normal); soft-preemption { cleanup-timer <i>seconds</i>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Control RSVP session preemption.
Default	normal
Options	<p>aggressive—Preempt RSVP sessions whenever bandwidth is insufficient to handle all sessions. A session is preempted whenever bandwidth is lowered or a new higher-priority session is established.</p> <p>disabled—Do not preempt RSVP sessions.</p> <p>normal—Preempt RSVP sessions to accommodate new higher-priority sessions when bandwidth is insufficient to handle all sessions.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Preempting RSVP Sessions on page 484

priority (Protocols RSVP)

Syntax	<code>priority setup-priority reservation-priority;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the setup priority and reservation priority for a bypass LSP. If insufficient link bandwidth is available during session establishment, the setup priority is compared with other setup priorities for established sessions on the link to determine whether some of them should be preempted to accommodate the new session. The session with the lower-hold priority is preempted.
Options	<p>reservation-priority—Reservation priority, used to keep a reservation after it has been set up. A smaller number has a higher priority. The priority must be greater than or equal to the setup priority to prevent preemption loops.</p> <p>Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p>Default: 0 (Once the session is set up, no other session can preempt it.)</p> <p>setup-priority—Setup priority.</p> <p>Range: 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p>Default: 7 (The session cannot preempt any existing sessions.)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Priority and Preemption for Bypass LSPs on page 509 • Configuring Priority and Preemption for LSPs on page 250

refresh-time

Syntax	refresh-time <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Set the refresh time.
Options	<i>seconds</i> —Refresh time. Range: 1 through 65,535 Default: 30 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Timers for RSVP Refresh Messages on page 483

reliable

Syntax	(reliable no-reliable);
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp peer-interface <i>peer-interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable reliable message delivery on the interface. In order to have refresh reduction and reliable delivery, you must include the aggregate and reliable statements.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RSVP Refresh Reduction on page 473• aggregate on page 974

rsvp

Syntax	rsvp;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.5 for EX Series switches.
Description	<p>Enable Resource Reservation Protocol (RSVP) signaling.</p> <p>You must include the rsvp statement in the configuration to enable RSVP on the router.</p> <p>The primary purpose of RSVP in Junos OS for EX Series switches is to support dynamic signaling within label switched paths (LSPs).</p>
Default	RSVP is disabled.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Minimum RSVP Configuration on page 471 • <i>Example: Configuring MPLS on EX8200 and EX4500 Switches</i> • <i>Configuring MPLS on Provider Edge EX8200 and EX4500 Switches Using Circuit Cross-Connect (CLI Procedure)</i> • <i>Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)</i> • <i>Configuring MPLS on EX8200 and EX4500 Provider Switches (CLI Procedure)</i>

rsvp-te (Routing Options)

Syntax	<pre>rsvp-te entry-name { destination-networks network-prefix; label-switched-path-template (Multicast) { default-template; template-name; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-options dynamic-tunnels <i>tunnel-name</i>], [edit routing-options dynamic-tunnels <i>tunnel-name</i>]
Release Information	Statement introduced in Junos OS Release 10.1. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Enable RSVP to automatically establish LSPs for any new PE router added to a full mesh of LSPs. To enable this feature, you must configure the rsvp-te statement on all of the PE routers in the full mesh.
Options	entry-name —Specify the entry for the RSVP tunnel. The other options are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RSVP Automatic Mesh on page 482• Flooding Unknown Traffic Using Point-to-Multipoint LSPs in VPLS

setup-protection

Syntax	setup-protection;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Description	The facility-backup fast reroute mechanism can provide setup protection for LSPs which are in the process of being signaled. Both point-to-point LSPs and point-to-multipoint LSPs are supported. You should configure the setup-protection statement on each of the routers along the LSP path on which you want to enable LSP setup protection. You should also configure IGP traffic engineering on all of the routers on the LSP path.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring RSVP Setup Protection on page 480

soft-preemption (Protocols RSVP)

Syntax	<code>soft-preemption { cleanup-timer <i>seconds</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp preemption], [edit protocols rsvp preemption]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable soft preemption to attempt to establish a new path for a preempted LSP before tearing it down.
Options	cleanup-timer —A value of 0 disables soft preemption. Range: 0 through 180 seconds Default: 30 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MPLS Soft Preemption on page 238

static-label-switched-path

```
Syntax  static-label-switched-path lsp-name {
        bypass bypass-name {
            bandwidth bps;
            description string;
            next-hop (address | interface-name | address/interface-name);
            push out-label;
            to address;
        }
        ingress {
            bandwidth bps;
            class-of-service cos-value;
            description string;
            install {
                destination-prefix <active>;
            }
            link-protection bypass-name name;
            metric metric;
            next-hop (address | interface-name | address/interface-name);
            node-protection bypass-name name next-next-label label;
            no-install-to-address;
            policing {
                filter filter-name;
                no-auto-policing;
            }
            preference preference;
            push out-label;
            to address;
        }
        transit incoming-label {
            bandwidth bps;
            description string;
            link-protection bypass-name name;
            next-hop (address | interface-name | address/interface-name);
            node-protection bypass-name name next-next-label label;
            pop;
            swap out-label;
        }
    }
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols mpls],
[edit protocols mpls]

Release Information Statement introduced in Junos OS Release 10.1.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Configure a static LSP.

Options *lsp-name*—Name of the path.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring Static LSPs on page 271](#)

subscription

Syntax

```
subscription percentage {
    ct0 percentage;
    ct1 percentage;
    ct2 percentage;
    ct3 percentage;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols rsvp interface *interface-name*],
[edit logical-systems *logical-system-name* protocols rsvp interface *interface-name* link-protection],
[edit protocols rsvp interface *interface-name*],
[edit protocols rsvp interface *interface-name* link-protection]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Configure the amount of bandwidth subscribed to a class type (when you have enabled Differentiated Services) or bypass LSP (when you have enabled link protection). **subscription** is the percentage of the link bandwidth that can be used for the RSVP reservation process.

Options **ctnumber percentage**—Percentage of the class-type bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type. You can specify bandwidth subscriptions for class types 0 through 3. This option is not available for bypass LSPs.

Range: 0 through 65,000

Default: 100 percent

percentage—Percentage of the class-type or bypass LSP bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type or bypass LSP.

Range: 0 through 65,000

Default: 100 percent

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Bandwidth Subscription Percentage for LSPs on page 323](#)
- [Configuring the Amount of Bandwidth Subscribed for Bypass LSPs on page 508](#)

traceoptions (Protocols RSVP)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable RSVP-level trace options.
Default	The default RSVP-level trace options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place RSVP tracing output in the file rsvp-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.</p> <ul style="list-style-type: none">• all—All tracing operations• error—All detected error conditions• event—RSVP-related events• lmp—RSVP-LMP interactions• packets—All RSVP packets• path—All path messages• pathtear—PathTear messages

- **resv**—Resv messages
- **resvtear**—ResvTear messages
- **route**—Routing information
- **state**—Session state transitions, including when RSVP-signaled LSPs come up and go down.

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-world-readable—(Optional) Enable only certain users to read the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Enable any user to read the log file.

Required Privilege	routing and trace—To view this statement in the configuration.
Level	routing-control and trace-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• Tracing RSVP Protocol Traffic on page 491
------------------------------	---

transit

Syntax	<pre>transit <i>incoming-label</i> { <i>bandwidth</i> <i>bps</i>; <i>description</i> <i>string</i>; <i>link-protection</i> <i>bypass-name name</i>; <i>next-hop</i> (<i>address</i> <i>interface-name</i> <i>address/interface-name</i>); <i>node-protection</i> <i>bypass-name name</i> <i>next-next-label label</i>; <i>pop</i>; <i>swap</i> <i>out-label</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i>], [edit protocols mpls static-label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in Junos OS Release 10.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series. Statement updated to include switch option in Junos OS Release 14.1X53-D25
Description	Configure a transit static LSP. The remaining statements are explained separately.
Options	<i>incoming-label</i> —Incoming label value. Range: 1000000 through 1048575
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Static LSPs on page 271

tunnel-services (RSVP)

Syntax	tunnel-services { devices <i>device-names</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable ultimate-hop popping on point-to-multipoint LSPs. The Junos OS selects one of the available virtual tunnel (VT) interfaces to de-encapsulate the egress traffic. By default, the selection process is performed automatically.
Default	Ultimate-hop popping is disabled.
Options	devices <i>device-names</i> —Specify which VT interfaces are used to handle the RSVP traffic. Range: 0 to 8 devices
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs on page 490

ultimate-hop-popping

Syntax	ultimate-hop-popping;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>label-switched-path-name</i>], [edit protocols mpls], [edit protocols mpls label-switched-path <i>label-switched-path-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3. Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.
Description	<p>Enable ultimate-hop popping on LSPs. Configure this statement on the device at the LSP ingress. In ultimate-hop popping, the MPLS label is popped from the IP packet at the PE router. The IP address is checked in a second address lookup (also at the PE router), and then the packet is forwarded to its destination.</p> <p>Be aware of the following platform requirements and restrictions:</p> <ul style="list-style-type: none"> • UHP LSPs using VT interfaces—Supported on all M Series, MX Series, T Series, and TX Matrix routers. • UHP LSPs using LSI interfaces—Supported on MX 3D Series routers only. • UHP LSP requirements for the egress PE device—For M Series and T Series routers, a VT interface is needed. • UHP LSPs and Layer 3 VPNs—UHP LSPs are supported for Layer 3 VPNs configured on MX 3D Series routers only. • UHP LSPs and VPLS—UHP LSPs are supported for VPLS configured on MX 3D Series routers only. You must configure the <i>no-tunnel-services</i> statement at the [edit routing-instances <i>routing-instance-name</i> protocols vpls] hierarchy level.
Default	Ultimate-hop popping is disabled by default on LSPs. Penultimate-hop popping is the default behavior. In penultimate-hop popping, the final MPLS label is popped from the IP packet at the last provider router in the network before being forwarded to the PE router. The PE router receives the packet and checks the IP address, and then the packet is forwarded to its destination.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Ultimate-Hop Popping for LSPs on page 222 • explicit-null on page 865

update-threshold

Syntax	update-threshold <i>threshold</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>], [edit protocols rsvp interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Adjust the threshold at which a change in bandwidth triggers an interior gateway protocol (IGP) update.
Options	threshold —Specify the percentage change in bandwidth to trigger an IGP update. Range: 1 through 20 percent Default: 10 percent
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the RSVP Update Threshold on an Interface on page 476

CHAPTER 25

LDP Configuration Statements

- [\[edit protocols ldp\] Hierarchy Level on page 1021](#)
- [allow-subnet-mismatch on page 1023](#)
- [authentication-algorithm on page 1024](#)
- [authentication-key \(Protocols LDP\) on page 1026](#)
- [authentication-key-chain \(Protocols LDP\) on page 1027](#)
- [auto-targeted-session on page 1028](#)
- [bfd-liveness-detection \(Protocols LDP\) on page 1029](#)
- [deaggregate on page 1030](#)
- [disable \(Protocols LDP\) on page 1031](#)
- [dod-request-policy on page 1032](#)
- [downstream-on-demand on page 1032](#)
- [ecmp on page 1033](#)
- [egress-policy on page 1033](#)
- [explicit-null \(Protocols LDP\) on page 1034](#)
- [export \(Protocols LDP\) on page 1034](#)
- [failure-action \(Protocols LDP\) on page 1035](#)
- [fec on page 1036](#)
- [graceful-restart \(Protocols LDP\) on page 1037](#)
- [hello-interval \(Protocols LDP\) on page 1038](#)
- [helper-disable \(LDP\) on page 1039](#)
- [holddown-interval on page 1040](#)
- [hold-time \(Protocols LDP\) on page 1041](#)
- [ignore-lsp-metrics on page 1042](#)
- [igp-synchronization on page 1042](#)
- [import \(Protocols LDP\) on page 1043](#)
- [ingress-policy on page 1044](#)
- [interface \(Protocols LDP\) on page 1045](#)
- [keepalive-interval on page 1046](#)

- [keepalive-timeout](#) on page 1047
- [l2-smart-policy](#) on page 1047
- [label-withdrawal-delay](#) on page 1048
- [ldp](#) on page 1049
- [ldp-synchronization](#) on page 1052
- [log-updown \(Protocols LDP\)](#) on page 1053
- [make-before-break \(LDP\)](#) on page 1054
- [maximum-neighbor-recovery-time](#) on page 1055
- [mldp-inband-signalling \(Protocols Multipoint LDP\)](#) on page 1056
- [mofrr-asm-starg \(Multicast-Only Fast Reroute in a PIM Domain\)](#) on page 1057
- [mofrr-disjoint-upstream-only \(Multicast-Only Fast Reroute in a PIM Domain\)](#) on page 1058
- [mofrr-no-backup-join \(Multicast-Only Fast Reroute in a PIM Domain\)](#) on page 1059
- [mofrr-primary-selection-by-routing \(Multicast-Only Fast Reroute\)](#) on page 1060
- [no-forwarding](#) on page 1061
- [oam \(Protocols LDP\)](#) on page 1062
- [p2mp \(Protocols LDP\)](#) on page 1063
- [p2mp-ldp-next-hop](#) on page 1064
- [periodic-traceroute](#) on page 1065
- [policing \(Protocols LDP\)](#) on page 1067
- [policy \(Multicast-Only Fast Reroute\)](#) on page 1068
- [policy \(Protocols Multipoint LDP\)](#) on page 1070
- [preference \(Protocols LDP\)](#) on page 1071
- [reconnect-time](#) on page 1072
- [recovery-time](#) on page 1073
- [session \(ldp\)](#) on page 1074
- [session-protection](#) on page 1075
- [stream-protection \(Multicast-Only Fast Reroute\)](#) on page 1076
- [strict-targeted-hellos](#) on page 1077
- [targeted-hello](#) on page 1077
- [traceoptions \(Protocols LDP\)](#) on page 1078
- [track-igp-metric](#) on page 1080
- [traffic-statistics \(Protocols LDP\)](#) on page 1081
- [transport-address](#) on page 1083
- [version \(BFD\)](#) on page 1084

[edit protocols ldp] Hierarchy Level

The following statements can also be configured at the **[edit logical-systems *logical-system-name*]** hierarchy level:

```

protocols {
  ldp {
    (deaggregate | no-deaggregate);
    egress-policy [ policy-names ];
    entropy-label {
      ingress-policy ingress-policy-name;
    }
    explicit-null;
    export [ policy-names ];
    graceful-restart {
      disable;
      helper-disable;
      maximum-neighbor-recovery-time seconds;
      reconnect-time seconds;
      recovery-time seconds;
    }
    import [ policy-names ];
    interface (interface-name | all) {
      disable;
      hello-interval seconds;
      hold-time seconds;
      transport-address (interface | router-id);
    }
    keepalive-interval seconds;
    keepalive-timeout seconds;
    log-updown {
      trap disable;
    }
    no-forwarding;
    oam {
      bfd-liveness-detection {
        detection-time threshold milliseconds;
        ecmp;
        failure-action {
          remove-nexthop;
          remove-route;
        }
        holddown-interval milliseconds;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        minimum-transmit-interval milliseconds;
        multiplier detection-time-multiplier;
        no-adaptation;
        transmit-interval {
          minimum-interval milliseconds;
          threshold milliseconds;
        }
      }
    }
    fec fec-address;
  }
}

```

```
ingress-policy ingress-policy-name;  
periodic-traceroute {  
    disable;  
    exp exp-value;  
    fanout fanout-value;  
    frequency minutes;  
    paths number-of-paths;  
    retries retry-attempts;  
    source address;  
    ttl ttl-value;  
    wait seconds;  
}  
}  
p2mp;  
root-address root-address{  
    lsp-id id;  
}  
policing {  
    fec fec-address {  
        ingress-traffic filter-name;  
        transit-traffic filter-name;  
    }  
}  
preference preference;  
session address {  
    authentication-key md5-authentication-key;  
}  
strict-targeted-hellos;  
traceoptions {  
    file filename <files number <size size> <world-readable | no-world-readable>;  
    flag flag <flag-modifier> <disable>;  
}  
track-igp-metric;  
traffic-statistics {  
    file filename <files number> <size size> <world-readable | no-world-readable>;  
    interval interval;  
    no-penultimate-hop;  
}  
transport-address (address | interface | router-id);  
}  
}
```

allow-subnet-mismatch

Syntax	allow-subnet-mismatch;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>], [edit protocols ldp interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Ignore the LDP subnet check. For Junos OS Release 8.4 and later releases, an LDP source address subnet check was added for the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address.
Default	The source address in the LDP link hello packet is matched against the interface address.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Ignoring the LDP Subnet Check on page 633

authentication-algorithm

Syntax authentication-algorithm *algorithm*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
 [edit logical-systems *logical-system-name* protocols bgp group *group-name*],
 [edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address*],
 [edit logical-systems *logical-system-name* protocols ldp session *session-address*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp session *session-address*],
 [edit logical-systems *logical-system-name* routing-options bmp],
 [edit logical-systems *logical-system-name* routing-options bmp station *station-name*],
 [edit protocols bgp],
 [edit protocols bgp group *group-name*],
 [edit protocols bgp group *group-name* neighbor *address*],
 [edit protocols ldp session *session-address*],
 [edit routing-instances *routing-instance-name* protocols bgp],
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*],
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*],
 [edit routing-instances *routing-instance-name* protocols ldp session *session-address*],
 [edit routing-options bmp],
 [edit routing-options bmp station *station-name*]

Release Information Statement introduced in Junos OS Release 7.6.
 Statement introduced for BGP in Junos OS Release 8.0.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
 Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.
 Statement introduced for BMP in Junos OS Release 13.3.
 Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure an authentication algorithm type.



NOTE: Keep the following points in mind when you configure the authentication algorithm in an IPsec proposal:

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, an error occurs and the tunnel is not established in this scenario. For example, if one end of the tunnel contains router 1 configured with the authentication algorithm as hmac-sha-256-128 and the other end of the tunnel contains router 2 configured with the authentication algorithm as hmac-md5-96, the VPN tunnel is not established.

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, and when one end of the tunnel contains two IPsec proposals to check whether a less secure algorithm is selected or not, an error occurs and the tunnel is not established. For example, if you configure two authentication algorithms for an IPsec proposal as `hmac-sha-256-128` and `hmac-md5-96` on one end of the tunnel, router 1, and if you configure the algorithm for an IPsec proposal as `hmac-md5-96` on the other end of the tunnel, router 2, the tunnel is not established and the number of proposals mismatch.
 - When you configure two IPsec proposals at both ends of a tunnel, such as the `authentication-algorithm hmac-sha-256-128` and `authentication-algorithm hmac-md5-96` statements at the `[edit services ipsec-vpn ipsec proposal proposal-name]` hierarchy level on one of the tunnel, router 1 (with the algorithms in two successive statements to specify the order), and the `authentication-algorithm hmac-md5-96` and `authentication-algorithm hmac-sha-256-128` statements at the `[edit services ipsec-vpn ipsec proposal proposal-name]` hierarchy level on one of the tunnel, router 2 (with the algorithms in two successive statements to specify the order, which is the reverse order of router 1), the tunnel is established in this combination as expected because the number of proposals is the same on both ends and they contain the same set of algorithms. However, the authentication algorithm selected is `hmac-md5-96` and not the stronger algorithm of `hmac-sha-256-128`. This method of selection of the algorithm occurs because the first matching proposal is selected. Also, for a default proposal, regardless of whether the router supports the Advanced Encryption Standard (AES) encryption algorithm, the `3des-cbc` algorithm is chosen and not the `aes-cfb` algorithm, which is because of the first algorithm in the default proposal being selected. In the sample scenario described here, on router 2, if you reverse the order of the algorithm configuration in the proposal so that it is the same order as the one specified on router 1, `hmac-sha-256-128` is selected as the authentication method.
 - You must be aware of the order of proposals in an IPsec policy at the time of configuration if you want the matching of proposals to happen in a certain order of preference, such as the strongest algorithm to be considered first when a match is made when both policies from the two peers have a proposal.
-

- Options** *algorithm*—Specify one of the following types of authentication algorithms:
- **aes-128-cmac-96**—Cipher-based message authentication code (AES128, 96 bits).
 - **hmac-sha-1-96**—Hash-based message authentication code (SHA1, 96 bits).
 - **md5**—Message digest 5.
- Default:** hmac-sha-1-96



NOTE: The default is not displayed in the output of the `show bgp bmp` command unless a key or key-chain is also configured.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring Route Authentication for BGP](#)
- [Configuring BGP Monitoring Protocol Version 3](#)

authentication-key (Protocols LDP)

Syntax authentication-key *md5-authentication-key*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp session *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp session *address*],
[edit protocols ldp session *address*],
[edit routing-instances *routing-instance-name* protocols ldp session *address*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Configure the MD5 authentication signature. The maximum length of the authentication signature is 69 characters.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring the TCP MD5 Signature for LDP Sessions on page 630](#)

authentication-key-chain (Protocols LDP)

Syntax	authentication-key-chain <i>key-chain</i> ;
Hierarchy Level	[edit logical-systems <i>name</i> protocols ldp session <i>address</i>], [edit logical-systems <i>name</i> routing-instances <i>instance-name</i> protocols ldp session <i>address</i>], [edit protocols ldp session <i>address</i>], [edit routing-instances <i>instance-name</i> protocols ldp session <i>address</i>]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for LDP, you cannot commit the 0.0.0.0/allow statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations.
Options	key-chain —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols • Configuring Miscellaneous LDP Properties on page 628

auto-targeted-session

Syntax	auto-targeted-session { maximum-sessions <i>seconds</i> ; teardown-delay <i>seconds</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit protocols ldp]
Release Information	Statement introduced in Junos OS Release 14.2.
Description	Configure session parameters for LDP sessions established with the remote LFA node that are automatically targeted using the loopback addresses. Configure parameters of automatically targeted sessions for remote LFA only.
Options	<p>maximum-sessions <i>seconds</i> —Specify the maximum number of auto-targeted LDP sessions allowed. Include this statement to optimize the use of router memory.</p> <p>Default: 100</p> <p>Range: 1 through 1000</p> <p>teardown-delay <i>seconds</i> —Specify the minimum time period for which the auto-targeted session must be alive before tearing down the auto-targeted LDP sessions to the remote LFA node. Include this statement to prevent rapid route-resolution in case of temporary change in IGP topology.</p> <p>Default: 90 seconds</p> <p>Range: 1 through 300 seconds</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>no-eligible-remote-backup</i>• <i>remote-backup-calculation</i>

bfd-liveness-detection (Protocols LDP)

Syntax	<pre> bfd-liveness-detection { detection-time threshold <i>milliseconds</i>; ecmp; failure-action { remove-nexthop; remove-route; } holddown-interval <i>seconds</i>; minimum-interval <i>milliseconds</i>; minimum-receive-interval <i>milliseconds</i>; minimum-transmit-interval <i>milliseconds</i>; multiplier <i>detection-time-multiplier</i>; no-adaptation; transmit-interval { minimum-interval <i>milliseconds</i>; threshold <i>milliseconds</i>; } version (0 1 automatic); } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp oam], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec address], [edit protocols ldp oam], [edit protocols ldp oam fec address]</p>
Release Information	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Support for the bfd-liveness-detection statement at the [edit protocols ldp oam fec address] hierarchy level and the ecmp option added in Junos OS Release 9.0.</p> <p>Support for the failure-action statement with the remove-nexthop and remove-route options and the holddown-interval statement added in Junos OS Release 9.4.</p>
Description	<p>Enable Bidirectional Forwarding Detection (BFD) for all MPLS LSPs or for just a specific LSP.</p>
Options	<p>minimum-interval—Minimum transmit and receive interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>minimum-receive-interval—Minimum receive interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>minimum-transmit-interval—Minimum transmit interval. Range: 50 through 255,000 milliseconds Default: 50</p> <p>multiplier—Detection time multiplier. Range: 50 through 255</p>

Default: 3

The other options are explained separately.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring BFD for LDP LSPs on page 543

deaggregate

Syntax	deaggregate no-deaggregate;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Control forwarding equivalence class (FEC) deaggregation on the router. The use of the deaggregate statement in LDP is a standard practice that we recommend for LDP deployments.
Default	Deaggregation is disabled on the router.
Options	deaggregate —Deaggregate FECs. no-deaggregate —Aggregate FECs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring FEC Deaggregation on page 541

disable (Protocols LDP)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options graceful-restart], [edit protocols ldp graceful-restart], [edit protocols ldp interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Explicitly disable LDP on an interface, or explicitly disable LDP graceful restart.
Default	LDP is enabled on interfaces configured with the LDP interface statement. LDP graceful restart is automatically enabled when graceful restart is enabled under the [edit routing-options] hierarchy level.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling and Disabling LDP on page 528 • Configuring LDP Graceful Restart on page 533

dod-request-policy

Syntax	<code>dod-request-policy <i>dod-request-policy-name</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit protocols ldp]
Release Information	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the name of the LDP downstream on demand request policy. LDP sends label request messages only for those FECs matching in the downstream on demand request policy.
Options	<i>dod-request-policy-name</i> —Specify the name of the downstream on demand request policy.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring LDP Downstream on Demand on page 592

downstream-on-demand

Syntax	<code>downstream-on-demand;</code>
Hierarchy Level	[edit logical systems <i>logical-system-name</i> protocols ldp session <i>session-address</i>], [edit protocols ldp session <i>session-address</i>]
Release Information	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable LDP downstream on demand on the LDP session. LDP is widely deployed in downstream unsolicited advertisement mode. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between access and aggregation networks to minimize the workload for the access node (AN) control plane and to avoid the storage of tens of thousands of label bindings from upstream aggregation nodes.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring LDP Downstream on Demand on page 592

ecmp

Syntax	<code>ecmp;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-liveness-detection], [edit protocols ldp oam bfd-liveness-detection], [edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Allows LDP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the ecmp statement, you must also configure the periodic-traceroute statement for the specified FEC. If you do not do so, the commit operation fails. You can configure the periodic-traceroute statement at the global hierarchy level ([edit protocols ldp oam]) while only configuring the ecmp statement for a specific FEC ([edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]).
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring ECMP-Aware BFD for LDP LSPs on page 546

egress-policy

Syntax	<code>egress-policy [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Control the prefixes advertised into LDP.
Default	Only the loopback address is advertised.
Options	<i>policy-names</i> —Name of one or more routing policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Prefixes Advertised into LDP from the Routing Table on page 540

explicit-null (Protocols LDP)

Syntax	explicit-null;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Advertise label 0 to the egress router of a label-switched path (LSP).
Default	If you do not include the explicit-null statement in the MPLS configuration, label 3 (implicit null) is advertised.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router on page 629

export (Protocols LDP)

Syntax	export [<i>policy-names</i>];
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Apply policy filters to outbound LDP label bindings. Filters are applied to all label bindings from all neighbors.
Options	<i>policy-names</i> —Name of one or more routing policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Filtering Outbound LDP Label Bindings on page 537

failure-action (Protocols LDP)

Syntax	<pre>failure-action { remove-nexthop; remove-route; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-liveness-detection],</p> <p>[edit protocols ldp oam bfd-liveness-detection],</p> <p>[edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]</p>
Release Information	Statement introduced in Junos OS Release 9.4.
Description	Configure route and next-hop properties in the event of a BFD session failure event on an LDP LSP. The failure event could be an existing BFD session that has gone down or could be a BFD session that never came up. LDP adds back the route or next hop when the relevant BFD session comes back up.
Options	<p>remove-nexthop—Remove a route corresponding to a next hop of the LSP's route at the ingress node when a BFD session failure event is detected.</p> <p>remove-route—Remove the route corresponding to an LSP from the appropriate routing tables when a BFD session failure event is detected. If the LSP is configured with ECMP and a BFD session corresponding to any path goes down, the route is removed.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a Failure Action for the BFD Session on an LDP LSP on page 546

fec

```

Syntax  fec fec-address {
        bfd-liveness-detection {
            detection-time threshold milliseconds;
            ecmp;
            failure-action {
                remove-nexthop;
                remove-route;
            }
            holddown-interval milliseconds;
            ingress-policy ingress-policy-name;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-transmit-interval milliseconds;
            multiplier detection-time-multiplier;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (0 | 1 | automatic);
        }
        no-bfd-liveness-detection;
        periodic-traceroute {
            disable;
            exp exp-value;
            fanout fanout-value;
            frequency minutes;
            paths number-of-paths;
            retries retry-attempts;
            source address;
            ttl ttl-value;
            wait seconds;
        }
    }

```

Hierarchy Level [edit logical-systems *logical-systems-name* protocols ldp oam],
[edit protocols ldp oam]

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 12.2 for EX Series switches.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Allows you to configure BFD for a specific LDP forwarding equivalence class (FEC).

Options *fec-address*—Specify the FEC address.

The other statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

- Related Documentation**
- [Configuring BFD for LDP LSPs on page 543](#)

graceful-restart (Protocols LDP)

Syntax	<pre>graceful-restart { disable; helper-disable; maximum-neighbor-recovery-time <i>value</i>; reconnect-time <i>seconds</i>; recovery-time <i>value</i>; }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	Configure LDP graceful restart on the LDP master protocol instance or for a specific routing instance.



NOTE: When you alter the graceful restart configuration at either the [edit routing-options graceful-restart] or [edit protocols ldp graceful-restart] hierarchy levels, any running LDP session is automatically restarted to apply the graceful restart configuration. This behavior mirrors the behavior of BGP when you alter its graceful restart configuration.

Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
---------------------------------	--

- Related Documentation**
- [Configuring LDP Graceful Restart on page 533](#)

hello-interval (Protocols LDP)

Syntax	<code>hello-interval <i>seconds</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> ldp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code> ldp targeted-hello],</code> <code>[edit protocols ldp interface <i>interface-name</i>],</code> <code>[edit protocols ldp targeted-hello],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Support for LDP targeted hellos added in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Control the LDP timer that regulates how often hello messages are sent. You can control the rate both link hello messages and targeted hello messages are sent depending on the hierarchy level at which you configure the hello-interval statement.
Options	<i>seconds</i> —Length of time between transmission of hello packets. Range: 1 through 65,535 seconds Default: 5 seconds for link hello messages, 15 seconds for targeted hello messages
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the LDP Timer for Hello Messages on page 528

helper-disable (LDP)

Syntax	helper-disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Disable helper mode for LDP graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart LDP.
Default	Helper mode is enabled by default on all routing protocols (including LDP) that support graceful restart.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring LDP Graceful Restart on page 533

holddown-interval

Syntax	<code>holddown-interval <i>holddown-interval</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-livenesss-detection], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-livenesss-detection], [edit protocols ldp oam bfd-livenesss-detection], [edit protocols ldp oam fec <i>address</i> bfd-livenesss-detection]
Release Information	Statement introduced in Junos OS Release 9.4.
Description	Specify how long the BFD session should be up before adding the route or next hop. Specifying a time of 0 seconds causes the route or next hop to be added as soon as the BFD session comes back up.
Options	<i>holddown-interval</i> —Number of seconds the BFD session should remain up before adding the route or next hop. Default: 0 seconds Range: 0 through 65,535 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Holddown Interval for the BFD Session on page 547

hold-time (Protocols LDP)

Syntax	<code>hold-time <i>seconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello],</p> <p>[edit protocols ldp interface <i>interface-name</i>],</p> <p>[edit protocols ldp targeted-hello],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support for LDP targeted hellos added in Junos OS Release 9.5.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	<p>Specify how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. You can specify times for both link hello messages and targeted hello messages depending on the hierarchy level at which you configure the hold-time statement.</p>
Options	<p><i>seconds</i>—Hold-time value.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 15 seconds for link hello messages, 45 seconds for targeted hello messages</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Delay Before LDP Neighbors Are Considered Down on page 529

ignore-lsp-metrics

Syntax	ignore-lsp-metrics;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf traffic-engineering shortcuts], [edit protocols ospf traffic-engineering shortcuts]
Release Information	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Cause OSPF to ignore the RSVP LSP metric. Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This can cause interoperability problems when you configure LDP tunneling over RSVP LSPs in heterogeneous networks.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks on page 630

igp-synchronization

Syntax	igp-synchronization holddown-interval <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.
Options	holddown-interval <i>seconds</i> —Time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. Default: 10 seconds Range: 10 through 60 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring LDP Synchronization with the IGP on the Router on page 633

import (Protocols LDP)

Syntax	<code>import [<i>policy-names</i>];</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Apply policy filters to received LDP label bindings. Filters are applied to all label bindings from all neighbors.
Options	<i>policy-names</i> —Name of one or more routing policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Filtering Inbound LDP Label Bindings on page 535

ingress-policy

Syntax	<code>ingress-policy [<i>ingress-policy-names</i>];</code>
Hierarchy Level	[edit logical-system <i>logical-system-name</i> protocols ldp entropy-label], [edit logical-system <i>logical-system-name</i> protocols ldp oam], [edit protocols ldp entropy-label], [edit protocols ldp oam]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced at the [edit protocols ldp entropy-label] hierarchy level in Junos OS Release 14.1.
Description	<p>Configure an LDP ingress policy for either the entropy label or Operation, Administration, and Management (OAM).</p> <p>For OAM, configure the ingress policy to choose which forwarding equivalence classes (FECs) need to have OAM enabled. If the FEC passes through the policy or if the FEC is explicitly configured, OAM is enabled for a FEC. For FECs chosen using a policy, the BFD parameters configured under [edit protocols ldp oam bfd-liveness-detection] are applied.</p>
Options	<i>ingress-policy-names</i> —Specify the names of the ingress policies.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring OAM Ingress Policies for LDP on page 547• Configuring the Entropy Label for LSPs on page 218

interface (Protocols LDP)

Syntax	<pre>interface <i>interface-name</i> { disable; hello-interval <i>seconds</i>; hold-time <i>seconds</i>; transport-address (interface loopback); }</pre>
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]</pre>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	Enable LDP on one or more router interfaces.
Default	LDP is disabled on all interfaces.
Options	<p><i>interface-name</i>—Name of an interface. To configure all interfaces, specify all.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Enabling and Disabling LDP on page 528

keepalive-interval

Syntax	<code>keepalive-interval <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Set the keepalive interval value.
Options	<i>seconds</i> —Keepalive value. Range: 1 through 65,535 Default: 10 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interval for LDP Keepalive Messages on page 531

keepalive-timeout

Syntax	<code>keepalive-timeout <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Set the keepalive timeout value. The keepalive timeout defines the amount of time that the neighbor LDP node waits before determining that the session has failed.
Options	<i>seconds</i> —Keepalive timeout value. Range: 1 through 65,535 Default: 30 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the LDP Keepalive Timeout on page 531

l2-smart-policy

Syntax	<code>l2-smart-policy;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Prevent LDP from exporting IPv4 FECs over sessions with Layer 2 neighbors only. IPv4 FECs received over such sessions are filtered out.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring LDP IPv4 FEC Filtering on page 542

label-withdrawal-delay

Syntax	label-withdrawal-delay <i>seconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Delay the withdrawal of labels to reduce router workload during IGP convergence.
Options	seconds —Configure the number of seconds to wait before withdrawing labels for the LDP LSPs. Default: 60 seconds Range: 0 through 300 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Label Withdrawal Timer on page 633

ldp

```
Syntax  ldp {
    (deaggregate | no-deaggregate);
    egress-policy [ policy-names ];
    explicit-null;
    export [ policy-names ];
    graceful-restart {
        disable;
        helper-disable;
        maximum-neighbor-recovery-time seconds;
        reconnect-time seconds;
        recovery-time seconds;
    }
    import [ policy-names ];
    interface (interface-name | all) {
        disable;
        hello-interval seconds;
        hold-time seconds;
        transport-address (interface | router-id);
    }
    keepalive-interval seconds;
    keepalive-timeout seconds;
    log-updown {
        trap disable;
    }
    no-forwarding;
    oam {
        bfd-liveness-detection {
            detection-time threshold milliseconds;
            ecmp;
            failure-action {
                remove-nexthop;
                remove-route;
            }
            holddown-interval milliseconds;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-transmit-interval milliseconds;
            multiplier detection-time-multiplier;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
        }
    }
    fec fec-address {
        bfd-liveness-detection {
            detection-time threshold milliseconds;
            ecmp;
            failure-action {
                remove-nexthop;
                remove-route;
            }
        }
    }
}
```

```

    holddown-interval milliseconds;
    ingress-policy ingress-policy-name;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-transmit-interval milliseconds;
    multiplier detection-time-multiplier;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (0 | 1 | automatic);
}
no-bfd-liveness-detection;
periodic-traceroute {
    disable;
    exp exp-value;
    fanout fanout-value;
    frequency minutes;
    paths number-of-paths;
    retries retry-attempts;
    source address;
    ttl ttl-value;
    wait seconds;
}
}
ingress-policy ingress-policy-name;
periodic-traceroute {
    disable;
    exp exp-value;
    fanout fanout-value;
    frequency minutes;
    paths number-of-paths;
    retries retry-attempts;
    source address;
    ttl ttl-value;
    wait seconds;
}
}
p2mp;
policing {
    fec fec-address {
        ingress-traffic filter-name;
        transit-traffic filter-name;
    }
}
}
preference preference;
session address {
    authentication-algorithm algorithm;
    authentication-key authentication-key;
    authentication-key-chain key-chain-name;
}
}
strict-targeted-hellos;
traceoptions {
    file filename <files number <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}

```



```

}
track-igp-metric;
traffic-statistics {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval interval;
  no-penultimate-hop;
}
transport-address (address | interface | router-id);
}

```

Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable LDP routing on the router or switch. You must include the ldp statement in the configuration to enable LDP on the router or switch.
Default	LDP is disabled on the router.
Options	The other statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Minimum LDP Configuration on page 528 • Enabling and Disabling LDP on page 528

ldp-synchronization

Syntax	<pre>ldp-synchronization { disable; hold-time seconds; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf interface <i>interface-name</i>], [edit protocols ospf interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ospf interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Enable synchronization by advertising the maximum cost metric until LDP is operational on the link.
Options	The other statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring LDP Synchronization with the IGP on LDP Links on page 632

log-updown (Protocols LDP)

Syntax	log-updown { trap disable; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Disable LDP traps on the router, logical system, or routing instance.
Options	trap disable —Disable LDP traps. Default: LDP traps are enabled on the router.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Disabling SNMP Traps for LDP on page 632

make-before-break (LDP)

Syntax	<pre>make-before-break { timeout <i>seconds</i>; switchover-delay <i>seconds</i>; }</pre>
Hierarchy Level	[edit protocols ldp]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	Configures make before break (MBB) for multicast LDP (MLDP) link protection to ensure minimum packet loss when attempting to signal a new label-switched path (LSP) before tearing down the old LSP path.
Options	<p>timeout <i>seconds</i>—Specify a value to change a make -before-break timeout for point-to-multipoint LSPs. Even if an MBB acknowledgment is not received for a point-to-multipoint LSP before the specified timeout period expires, the label-switching router (LSR) performs an MBB switchover from the old LSR to the new upstream LSR.</p> <p>Range: 1 through 300 seconds</p> <p>Default: 30 seconds</p> <p>switchover-delay <i>seconds</i>—Specify a value to change switchover delay for a point-to-multipoint LSP from the old LSR to the new upstream LSR. If an MBB acknowledgment is received on a point of local repair (PLR) router, the PLR waits for the specified seconds to switch its upstream LSR from the old LSR to the new LSR.</p> <p>Range: 1 through 300 seconds</p> <p>Default: 30 seconds</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring LDP Link Protection on page 549

maximum-neighbor-recovery-time

Syntax	<code>maximum-neighbor-recovery-time <i>seconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement changed from maximum-recovery-time to maximum-neighbor-recovery-time in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.
Options	<i>seconds</i> —Configure the maximum recovery time, in seconds. Range: 120 through 1800 seconds Default: 140 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Recovery Time and Maximum Recovery Time on page 535 • Configuring Graceful Restart Options for LDP • no-strict-lsa-checking • recovery-time

mldp-inband-signalling (Protocols Multipoint LDP)

Syntax	mldp-inband-signalling { policy <i>policy-name</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols pim], [edit protocols pim],
Release Information	Statement introduced in Junos OS Release 13.2.
Description	<p>Multipoint LDP (M-LDP) in-band signaling enables you to carry multicast traffic across an existing IP/MPLS backbone, while avoiding the use of PIM in the provider core.</p> <p>On the label-edge router (LER), enable PIM to use M-LDP in-band signaling for the upstream neighbors when the LER does not detect a PIM upstream neighbor. On the egress nodes, configure the MPLS LSP root in the PIM configuration, using the policy statement.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs on page 597

mofrr-asm-starg (Multicast-Only Fast Reroute in a PIM Domain)

Syntax	<code>mofrr-asm-starg;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast stream-protection], [edit logical-systems <i>logical-system-name</i> routing-options multicast stream-protection], [edit routing-instances <i>routing-instance-name</i> routing-options multicast stream-protection], [edit routing-options multicast stream-protection]
Release Information	Statement introduced in Junos OS Release 14.1.
Description	Enable mofrr-asm-starg to include any-source multicast (ASM) for (*G) joins in the Multicast-only fast reroute (MoFRR).



NOTE: **mofrr-asm-starg** applies to IP-PIM only. When enabled for group G, *G will undergo MoFRR as long as there is no S#,G for Group G. In other words, *G MoFRR will cease and any old states will be torn down when S#,G is created. Note too, that **mofrr-asm-starg** is not supported for mLDP (since mLDP itself does not support *G).

In a PIM domain with MoFRR enabled, the default for **stream-protection** is S,G routes only.

Context: Multicast-only fast reroute (MoFRR) can be used to reduce traffic loss in a multicast distribution tree in the event of link down. To employ MoFRR, a downstream router is configured with an alternative path back towards the source, over which it receives a backup live stream of the same multicast traffic. That router propagates the same (S,G) join toward both upstream neighbors in order to create duplicate multicast trees. If a failure is detected on the primary tree, the router switches to the backup tree to prevent packet loss.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Multicast-Only Fast Reroute on page 566 • Example: Configuring Multicast-Only Fast Reroute in a PIM Domain • Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain on page 576

mofrr-disjoint-upstream-only (Multicast-Only Fast Reroute in a PIM Domain)

Syntax	<code>mofrr-disjoint-upstream-only;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast stream-protection],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options multicast stream-protection],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast stream-protection],</code> <code>[edit routing-options multicast stream-protection]</code>
Release Information	Statement introduced in Junos OS Release 14.1.
Description	<p>When you configure multicast-only fast reroute (MoFRR) in a PIM domain, allow only a disjoint RPF (an RPF on a separate plane) to be selected as the backup RPF path.</p> <p>In a multipoint LDP MoFRR domain, the same label is shared between parallel links to the same upstream neighbor. This is not the case in a PIM domain, where each link forms a neighbor. The mofrr-disjoint-upstream-only statement does not allow a backup RPF path to be selected if the path goes to the same upstream neighbor as that of the primary RPF path. This ensures that MoFRR is triggered only on a topology that has multiple RPF upstream neighbors.</p>
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Multicast-Only Fast Reroute on page 566• Example: Configuring Multicast-Only Fast Reroute in a PIM Domain• Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain on page 576

mofrr-no-backup-join (Multicast-Only Fast Reroute in a PIM Domain)

Syntax	mofrr-no-backup-join;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast stream-protection], [edit logical-systems <i>logical-system-name</i> routing-options multicast stream-protection], [edit routing-instances <i>routing-instance-name</i> routing-options multicast stream-protection], [edit routing-options multicast stream-protection]
Release Information	Statement introduced in Junos OS Release 14.1.
Description	When you configure multicast-only fast reroute (MoFRR) in a PIM domain, prevent sending join messages on the backup path, but retain all other MoFRR functionality.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Multicast-Only Fast Reroute on page 566 • Example: Configuring Multicast-Only Fast Reroute in a PIM Domain • Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain on page 576

mofrr-primary-selection-by-routing (Multicast-Only Fast Reroute)

Syntax	<code>mofrr-primary-selection-by-routing;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast stream-protection],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-options multicast stream-protection],</code> <code>[edit routing-instances <i>routing-instance-name</i> routing-options multicast stream-protection],</code> <code>[edit routing-options multicast stream-protection]</code>
Release Information	Statement introduced in Junos OS Release 14.1.
Description	<p>When you configure multicast-only fast reroute (MoFRR) in a PIM domain, allow new primary path selection to be based on the unicast gateway selection for the unicast route to the source and to change when there is a change in the unicast selection, rather than having the backup path be promoted as primary. This ensures that the primary RPF hop is always on the best path.</p> <p>When you include the mofrr-primary-selection-by-routing statement, the backup path is not guaranteed to get promoted to be the new primary path when the primary path goes down.</p>
Default	By default, the backup path gets promoted to be the primary path when MoFRR is configured in a PIM domain.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Multicast-Only Fast Reroute on page 566• Example: Configuring Multicast-Only Fast Reroute in a PIM Domain• Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain on page 576

no-forwarding

Syntax	no-forwarding;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Do not add ingress routes to the inet.0 routing table even if traffic-engineering bgp-igp (configured at the [edit protocols mpls] hierarchy level) is enabled.
Default	The no-forwarding statement is disabled. Ingress routes are added to the inet.0 routing table instead of the inet.3 routing table when traffic-engineering bgp-igp is enabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Preventing Addition of Ingress Routes to the inet.0 Routing Table on page 628 • Configuring Virtual-Router Routing Instances in VPNs

oam (Protocols LDP)

```
Syntax  oam {
        bfd-liveness-detection {
            detection-time threshold milliseconds;
            ecmp;
            failure-action {
                remove-nexthop;
                remove-route;
            }
            holddown-interval milliseconds;
            ingress-policy ingress-policy-name;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-transmit-interval milliseconds;
            multiplier detection-time-multiplier;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
            version (0 | 1 | automatic);
        }
        fec fec-address;
        ingress-policy ingress-policy-name;
        lsp-ping-interval seconds;
        periodic-traceroute {
            disable;
            exp exp-value;
            fanout fanout-value;
            frequency minutes;
            paths number-of-paths;
            retries retry-attempts;
            source address;
            ttl ttl-value;
            wait seconds;
        }
    }
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp]
[edit protocols ldp]

Release Information Statement introduced in Junos OS Release 7.6.
lsp-ping-interval option introduced in Junos OS Release 9.4.

Description Configure Operation, Administration, and Maintenance (OAM) and Bidirectional Forwarding Detection (BFD) protocol for LDP.

Options **fec *fec-address***—Specify the forwarding equivalence class (FEC) address. You must either specify a FEC address or configure an OAM ingress policy to ensure that the BFD session comes up.

lsp-ping-interval *seconds*—Specify the duration of the LSP ping interval in seconds. To issue a ping on an LDP-signaled LSP, use the **ping mpls ldp** command.

Default: 60 seconds

Range: 30 through 3,600 seconds

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring BFD for LDP LSPs on page 543](#)

p2mp (Protocols LDP)

Syntax p2mp{
 root-address *root-address*{
 lsp-id *id*;
 }

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced in Junos OS Release 11.2.

Description Enable point-to-multipoint MPLS LSPs in an LDP-signaled LSP.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs](#)
- [Point-to-Multipoint LSPs Overview on page 281](#)

p2mp-ldp-next-hop

Syntax	<pre>p2mp-ldp-next-hop { root-address <i>root-address</i>{ lsp-id <i>id</i>; } }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options static route <i>destination-prefix</i>], [edit logical-systems <i>logical-system-name</i> routing-options static route <i>destination-prefix</i>], [edit routing-instances <i>routing-instance-name</i> routing-options static route <i>destination-prefix</i>], [edit routing-options static route <i>destination-prefix</i>]
Release Information	Statement introduced in Junos OS Release 13.3.
Description	Specify a point-to-multipoint LDP label-switched path (LSP) as the next hop for a static route, and configure a root and provide an lsp-id on that LDP-signalled label-switched path.
Options	root-address <i>root address</i> — Specify the root address of the point-to-multipoint LSP. lsp-id <i>id</i> — Specify the generic LSP identifier. The range is 1 through 65535.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• [edit routing-options] <i>Hierarchy Level</i>

periodic-traceroute

Syntax	<pre> periodic-traceroute { disable; exp <i>exp-value</i>; fanout <i>fanout-value</i>; frequency <i>minutes</i>; paths <i>number-of-paths</i>; retries <i>retry-attempts</i>; source <i>address</i>; ttl <i>ttl-value</i>; wait <i>seconds</i>; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp oam], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>fec-address</i>], [edit protocols ldp oam], [edit protocols ldp oam fec <i>fec-address</i>]
Release Information	Statement introduced in Junos OS Release 8.4. Support added at the [edit protocols ldp oam] and [edit logical-systems <i>logical-system-name</i> protocols ldp oam] hierarchy levels in Junos OS Release 9.0. Statement introduced in Junos OS Release 12.2 for EX Series switches.
Description	Enable tracing of forwarding equivalence classes (FECs) for LDP LSPs.
Options	<p>disable—(Optional) Disable tracing for a specific FEC. This option is available at the [edit protocols ldp oam fec <i>fec-address</i> periodic-traceroute] and [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>fec-address</i> periodic-traceroute] hierarchy levels only.</p> <p>exp <i>exp-value</i>—(Optional) Specify the class of service to use when sending probes. Default: 7 Range: 0 through 7</p> <p>fanout <i>fanout-value</i>—(Optional) Specify the maximum number of next hops to search per node. Default: 16 Range: 1 through 16</p> <p>frequency <i>minutes</i>—(Optional) Specify the interval between traceroute attempts. Default: 60 minutes Range: 15 through 120 minutes</p> <p>paths <i>number-of-paths</i>—(Optional) Specify the maximum number of paths to search. Default: 3 Range: 1 through 255</p>

retries *retry-attempts*—(Optional) Specify the number of attempts to send a probe to a specific node before giving up.

Default: 3

Range: 1 through 9

source address—(Optional) Specify the IPv4 source address to use when sending probes.

ttl value—(Optional) Specify the maximum time-to-live value. Nodes that are beyond this value are not traced.

Default: 64

Range: 1 through 255

wait seconds—(Optional) Specify the wait interval before resending a probe packet.

Default: 10 seconds

Range: 5 though 15 seconds

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Configuring LDP LSP Traceroute on page 634
------------------------------	--

policing (Protocols LDP)

Syntax	<pre> policing { fec <i>fec-address</i> { ingress-traffic <i>filter-name</i>; transit-traffic <i>filter-name</i>; } } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],</p> <p>[edit protocols ldp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ldp]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	Enable policing of forwarding equivalence classes (FECs) for LDP.
Options	<p>fec <i>fec-address</i>—Specify the address for the FEC.</p> <p>ingress-traffic <i>filter-name</i>—Specify the name of the filter for policing ingress FEC traffic.</p> <p>transit-traffic <i>filter-name</i>—Specify the name of the filter for policing transit FEC traffic.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Policers for LDP FECs on page 541

policy (Multicast-Only Fast Reroute)

Syntax	<code>policy <i>policy-name</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast stream-protection], [edit logical-systems <i>logical-system-name</i> routing-options multicast stream-protection], [edit routing-instances <i>routing-instance-name</i> routing-options multicast stream-protection], [edit routing-options multicast stream-protection]
Release Information	Statement introduced in Junos OS Release 14.1.
Description	<p>When you configure multicast-only fast reroute (MoFRR), apply a routing policy that filters for a restricted set of multicast streams to be affected by your MoFRR configuration. You can apply filters that are based on source or group addresses.</p>

For example:

```

routing-options {
  multicast {
    stream-protection {
      policy mofrr-select;
    }
  }
}
policy-statement mofrr-select {
  term A {
    from {
      source-address-filter 225.1.1.1/32 exact;;
    }
    then {
      accept;
    }
  }
  term B {
    from {
      source-address-filter 226.0.0.0/8 orlonger;
    }
    then {
      accept;
    }
  }
  term C {
    from {
      source-address-filter 227.1.1.0/24 orlonger;
      source-address-filter 227.4.1.0/24 orlonger;
      source-address-filter 227.16.1.0/24 orlonger;
    }
    then {
      accept;
    }
  }
  term D {
    from {

```

```

        source-address-filter 227.1.1.1/32 exact
    }
    then {
        reject; #MoFRR disabled
    }
}
term E {
    from {
        route-filter 227.1.1.0/24 orlonger;
    }
    then accept;
}
...
}

```

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation

- [Understanding Multicast-Only Fast Reroute on page 566](#)
- [Example: Configuring Multicast-Only Fast Reroute in a PIM Domain](#)
- [Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain on page 576](#)

policy (Protocols Multipoint LDP)

Syntax	<code>policy <i>policy-name</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols pim mldp-inband-signalling], [edit protocols pim mldp-inband-signalling]
Release Information	Statement introduced in Junos OS Release 13.2.
Description	<p>Multipoint LDP (M-LDP) in-band signaling enables you to carry multicast traffic across an existing IP/MPLS backbone, while avoiding the use of PIM in the provider core.</p> <p>On the egress nodes of the point-to-multipoint LSP, specify an M-LDP join translation filter policy where PIM messages are translated into M-LDP FEC bindings. The policy statement is needed when internal BGP (IBGP) is not available in the core site or to override IBGP-based LSP root detection.</p> <p>The filter policy is configured at the [edit policy-options] hierarchy level. The policy generally specifies one or more source-address filters and the point-to-multipoint LDP root IP address using the p2mp-lsp-root policy action.</p>
Options	<i>policy-name</i> —Name of a policy configured at the [edit policy-options] hierarchy level.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs on page 597

preference (Protocols LDP)

Syntax	<code>preference <i>preference</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit protocols ldp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit protocols ldp interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Set the route preference level for LDP routes.
Options	<i>preference</i> —Preferred value. Range: 0 through 255 Default: 9
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring LDP Route Preferences on page 532

reconnect-time

Syntax	<code>reconnect-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.
Options	seconds —Time required for reconnection. Range: 30 through 300 Default: 60 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring LDP Graceful Restart on page 533 on <i>MPLS Applications Feature Guide for Routing Devices</i>• <i>Configuring Graceful Restart Options for LDP</i>

recovery-time

Syntax	<code>recovery-time seconds;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the amount of time a router waits for LDP to restart gracefully.
Options	seconds —Configure the recovery time, in seconds. Range: 120 through 1800 seconds Default: 140 seconds
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Recovery Time and Maximum Recovery Time on page 535

session (ldp)

Syntax	<pre>session address { authentication-algorithm <i>algorithm</i>; authentication-key <i>authentication-key</i>; authentication-key-chain <i>key-chain-name</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. authentication-algorithm statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the address for the remote end of the LDP session. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the TCP MD5 Signature for LDP Sessions on page 630

session-protection

Syntax	session-protection { timeout <i>seconds</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Description	Configure when an LDP session is torn down and resigaled after the router stops receiving hello messages from a neighboring router. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished. The LDP session remains up for the duration specified as long as the routers maintain IP network connectivity.
Options	timeout <i>seconds</i> —Time in seconds before the LDP session is torn down and resigaled. Range: 1 through 65,535 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring LDP Session Protection on page 631

stream-protection (Multicast-Only Fast Reroute)

Syntax	<pre>stream-protection { mofrr-asm-starg; mofrr-disjoint-upstream-only; mofrr-no-backup-join; mofrr-primary-selection-by-routing; policy <i>policy-name</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast], [edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-instances <i>routing-instance-name</i> routing-options multicast], [edit routing-options multicast]
Release Information	Statement introduced in Junos OS Release 14.1.
Description	Enable multicast-only fast reroute (MoFRR) on a routing device. MoFRR minimizes packet loss in a network when there is a link failure.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Multicast-Only Fast Reroute on page 566• Example: Configuring Multicast-Only Fast Reroute in a PIM Domain• Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain on page 576

strict-targeted-hellos

Syntax	strict-targeted-hellos;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Prevent LDP sessions from being established with remote neighbors that have not been specifically configured. LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling Strict Targeted Hello Messages for LDP on page 531

targeted-hello

Syntax	targeted-hello { hello-interval <i>seconds</i> ; hold-time <i>seconds</i> ; }
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
Release Information	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Specify the LDP timer and LDP hold time for targeted hellos.
Options	The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the LDP Timer for Hello Messages on page 528 • Configuring the Delay Before LDP Neighbors Are Considered Down on page 529

traceoptions (Protocols LDP)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; }</pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>match-on address option for the filter flag modifier added in Junos OS Release 10.4.</p> <p>nsr-synchronization and p2mp-nsr-synchronization operations for flag statement introduced in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	Specify LDP protocol-level trace options.
Default	The default LDP protocol-level trace options are inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>filename</i>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory ldp-log. We recommend that you place LDP tracing output in the file ldp-log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.</p> <ul style="list-style-type: none"> • address—Operation of address and address withdrawal messages • binding—Label-binding operations • error—Error conditions • event—Protocol events

- **initialization**—Operation of initialization messages
- **label**—Operation of label request, label map, label withdrawal, and label release messages
- **notification**—Operation of notification messages
- **nsr-synchronization**—Nonstop active routing synchronization events
- **p2mp-nsr-synchronization**—Point-to-multipoint nonstop active routing synchronization events
- **packets**—Equivalent to setting **address**, **initialization**, **label**, **notification**, and **periodic** flags (see also the **filter** flag modifier)
- **path**—Label-switched path operations
- **periodic**—Operation of hello and keepalive messages
- **route**—Operation of route messages
- **state**—Protocol state transitions

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **disable**—Disable this trace flag.
- **filter**—Filter to apply to this flag. The **filter** flag modifier can be applied only to the **route**, **path**, and **binding** flags. This flag modifier has the following options:
 - **match-on**—Match on argument specified. The **match-on** option has the following suboptions:
 - **address**—Filter based on the source and destination addresses of packets. Available for the **packets** flag option only.
 - **fec**—Filter based on the FEC associated with the traced object.
 - **policy** *policy-name*—Specify the filter policy.
 - **receive**—Packets being received.
 - **send**—Packets being transmitted.

no-world-readable—(Optional) Prevent all users from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Enable any user to read the log file.

Required Privilege Level routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- [Tracing LDP Protocol Traffic on page 637](#)
- *Network Management Administration Guide for Routing Devices*

track-igp-metric

Syntax track-igp-metric;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp],
[edit protocols ldp],
[edit routing-instances *routing-instance-name* protocols ldp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Cause the IGP route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring LDP to Use the IGP Route Metric on page 628](#)

traffic-statistics (Protocols LDP)

Syntax	<pre> traffic-statistics { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; interval <i>seconds</i>; no-penultimate-hop; } </pre>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],</p> <p>[edit protocols ldp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ldp]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	LDP traffic statistics display the amount of traffic passed through a router for a particular FEC.
Options	<p>file <i>filename</i>—Name of the file to receive the output of the LDP statistics operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log</code>.</p> <p>files <i>number</i>—(Optional) Maximum number of LDP statistics files. When a statistics file named <i>ldp-stat</i> reaches its maximum size, it is renamed <i>ldp-stat.0</i>, then <i>ldp-stat.1</i>, and so on, until the maximum number of LDP statistics files is reached. Then the oldest file is overwritten.</p> <p>Range: 2 through 1000</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you also must include the size statement to specify the maximum file size.</p> <p>interval <i>seconds</i>—(Optional) Specify the interval at which the statistics are polled and written to the file.</p> <p>Default: 300 seconds (5 minutes)</p> <p>no-penultimate-hop—(Optional) Do not collect traffic statistics on the penultimate hop router.</p> <p>no-world-readable—(Optional) Prevent all users from reading the log file.</p> <p>size <i>size</i>—(Optional) Maximum size of each statistics file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a statistics file named <i>ldp-stat</i> reaches this size, it is renamed <i>ldp-stat.0</i>. When <i>ldp-stat</i> again reaches this size, <i>ldp-stat.0</i> is renamed <i>ldp-stat.1</i> and <i>ldp-stat</i> is renamed <i>ldp-stat.0</i>. This renaming scheme continues until the maximum number of statistics files is reached. Then the oldest statistics file is overwritten.</p> <p>Syntax: <i>xk</i> to specify KB, <i>xm</i> to specify MB, or <i>xg</i> to specify GB</p> <p>Range: 10 KB through the maximum file size supported on your system</p>

Default: 1 MB

If you specify a maximum file size, you also must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Enable log file access for all users.

Required Privilege	routing—To view this statement in the configuration.
Level	routing-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• Collecting LDP Statistics on page 635
------------------------------	---

transport-address

Syntax	<code>transport-address (interface router-id);</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit protocols ldp interface <i>interface-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	<p>Enables you to configure the IP address used to specify the TCP session for the LDP session. Routers must first establish a TCP session between one another before they can establish an LDP session. The TCP session enables the routers to exchange the label advertisements needed for the LDP session. To establish the TCP session, each router must learn the other router's transport address. The transport address is an IP address used to identify the TCP session over which the LDP session will run.</p>
Default	router-id
Options	<p>interface—The first IP address on the interface is used as the transport address for any LDP sessions to neighbors that can be reached over that interface. You cannot specify the interface option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the router-id option.</p> <p>router-id—The router identifier is used as the transport address. Unless otherwise configured, the router identifier is the loopback address.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Specifying the Transport Address Used by LDP on page 539

version (BFD)

Syntax	version (0 1 automatic);
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-liveness-detection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ldp oam fec address bfd-liveness-detection],</p> <p>[edit system services dhcp-local-server liveness-detection method bfd],</p> <p>[edit system services dhcp-local-server dhcpv6 liveness-detection method bfd],</p> <p>[edit forwarding-options dhcp-relay liveness-detection method bfd],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method bfd],</p> <p>[edit system services dhcp-local-server group <i>group-name</i> liveness-detection method bfd],</p> <p>[edit system services dhcp-local-server dhcpv6 group <i>group-name</i> liveness-detection method bfd],</p> <p>[edit forwarding-options dhcp-relay group <i>group-name</i> liveness-detection method bfd],</p> <p>[edit forwarding-options dhcp-relay dhcpv6 group <i>group-name</i> liveness-detection method bfd],</p> <p>[edit protocols ldp oam bfd-liveness-detection],</p> <p>[edit protocols ldp oam fec address bfd-liveness-detection]</p>
Release Information	<p>Statement introduced in Junos OS Release 12.1.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p>
Description	Configure the BFD protocol version to detect.
Options	<p>0—Use BFD protocol version 0.</p> <p>1—Use BFD protocol version 1.</p> <p>automatic—Autodetect the BFD protocol version.</p> <p>Default: automatic</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Group Liveness Detection for DHCP Local Server Clients</i> • <i>Example: Configuring Global Liveness Detection for DHCP Relay Agent Clients</i> • Configuring BFD for LDP LSPs on page 543

CHAPTER 26

CCC and TCC Configuration Statements

- [connections \(Circuits\) on page 1086](#)
- [encapsulation \(Logical Interface\) on page 1087](#)
- [encapsulation \(Physical Interface\) on page 1091](#)
- [interface-switch on page 1096](#)
- [lsp-switch on page 1097](#)
- [output-interface \(CCC\) on page 1097](#)
- [p2mp-receive-switch on page 1098](#)
- [p2mp-transmit-switch on page 1099](#)
- [remote-interface-switch on page 1100](#)

connections (Circuits)

Syntax

```
connections {
  interface-switch connection-name {
    interface interface-name.unit-number;
  }
  lsp-switch connection-name {
    transmit-lsp label-switched-path;
    receive-lsp label-switched-path;
  }
  p2mp-receive-switch {
    output-interface [ interface-name.unit-number ];
    receive-p2mp-lsp receiving-point-to-multipoint-lsp;
  }
  p2mp-transmit-switch {
    input-interface interface-name.unit-number;
    transmit-p2mp-lsp transmitting-point-to-multipoint-lsp;
  }
  remote-interface-switch connection-name {
    interface interface-name.unit-number;
    receive-lsp label-switched-path;
    transmit-lsp label-switched-path;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
[edit protocols]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the connection between two circuits in a CCC connection.

Options The statements are explained separately.



NOTE: The edit logical-systems hierarchy is not available on QFabric systems.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [Configuring Layer 2 Switching Cross-Connects Using CCC on page 647](#)
- [Configuring MPLS LSP Tunnel Cross-Connects Using CCC on page 655](#)
- [Configuring TCC on page 659](#)
- [Configuring CCC Switching for Point-to-Multipoint LSPs on page 665](#)

encapsulation (Logical Interface)

Syntax	encapsulation (atm-ccc-cell-relay atm-ccc-vc-mux atm-cisco-nlpid atm-mlppp-llc atm-nlpid atm-ppp-llc atm-ppp-vc-mux atm-snap atm-tcc-snap atm-tcc-vc-mux atm-vc-mux ether-over-atm-llc ether-vpls-over-atm-llc ether-vpls-over-fr ether-vpls-over-ppp ethernet ethernet-ccc ethernet-vpls ethernet-vpls-fr frame-relay-ccc frame-relay-ether-type frame-relay-ether-type-tcc frame-relay-ppp frame-relay-tcc gre-fragmentation multilink-frame-relay-end-to-end multilink-ppp ppp-over-ether ppp-over-ether-over-atm-llc vlan-bridge vlan-ccc vlan-vci-ccc vlan-tcc vlan-vpls vxlan);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces rlsq <i>number</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers (ethernet , vlan-ccc , and vlan-tcc options only). Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access Routers. Only the atm-ccc-cell-relay and atm-ccc-vc-mux options are supported on ACX Series routers.
Description	Configure a logical link-layer encapsulation type.
Options	<p>atm-ccc-cell-relay—Use ATM cell-relay encapsulation.</p> <p>atm-ccc-vc-mux—Use ATM virtual circuit (VC) multiplex encapsulation on CCC circuits. When you use this encapsulation type, you can configure the ccc family only.</p> <p>atm-cisco-nlpid—Use Cisco ATM network layer protocol identifier (NLPID) encapsulation. When you use this encapsulation type, you can configure the inet family only.</p> <p>atm-mlppp-llc—For ATM2 IQ interfaces only, use Multilink Point-to-Point (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.</p> <p>atm-nlpid—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the inet family only.</p> <p>atm-ppp-llc—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over AAL5 LLC encapsulation.</p> <p>atm-ppp-vc-mux—(ATM2 IQ interfaces and MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP only) Use PPP over ATM AAL5 multiplex encapsulation.</p> <p>atm-snap—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM subnetwork attachment point (SNAP) encapsulation.</p> <p>atm-tcc-snap—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.</p>

atm-tcc-vc-mux—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

atm-vc-mux—(All interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the **inet** family only.

ether-over-atm-llc—(All IP interfaces including MX Series routers with MPC/MIC interfaces using the ATM MIC with SFP) For interfaces that carry IP traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.

ether-vpls-over-atm-llc—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

ether-vpls-over-fr—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, per RFC 2427, *Multiprotocol Interconnect over Frame Relay*.



NOTE: The SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP, and the DS3/E3 MIC do not support Ethernet over Frame Relay encapsulation.

ether-vpls-over-ppp—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Point-to-Point Protocol (PPP) encapsulation to support Bridged Ethernet over PPP-encapsulated TDM interfaces for VPLS applications.

ethernet—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.



NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE device over a time-division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ppp—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

gre-fragmentation—For adaptive services interfaces only, use GRE fragmentation encapsulation to enable fragmentation of IPv4 packets in GRE tunnels. This encapsulation clears the do not fragment (DF) bit in the packet header. If the packet's size exceeds the tunnel's maximum transmission unit (MTU) value, the packet is fragmented before encapsulation.

multilink-frame-relay-end-to-end—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

multilink-ppp—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

ppp-over-ether—You use PPP over Ethernet encapsulation to configure an underlying Ethernet interface for a dynamic PPPoE logical interface.

ppp-over-ether-over-atm-llc—(MX Series routers with MPCs using the ATM MIC with SFP only) For underlying ATM interfaces, use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

vlan-bridge—Use Ethernet VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q tagging, flexible-ethernet-services, and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

vlan-ccc—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-tcc—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-vpls—Use Ethernet VLAN encapsulation on VPLS circuits.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
---------------------------------	---

- | | |
|------------------------------|--|
| Related Documentation | <ul style="list-style-type: none">• Configuring Layer 2 Switching Cross-Connects Using CCC on page 647• Configuring the Encapsulation for Layer 2 Switching TCCs on page 659• <i>Configuring Interface Encapsulation on Logical Interfaces</i>• Configuring the CCC Encapsulation for LSP Tunnel Cross-Connects on page 656• <i>Circuit and Translational Cross-Connects Overview</i>• <i>Identifying the Access Concentrator</i>• <i>Configuring ATM Interface Encapsulation</i>• <i>Configuring VLAN Encapsulation</i>• <i>Configuring Extended VLAN Encapsulation</i>• <i>Configuring ATM-to-Ethernet Interworking</i>• <i>Configuring Interface Encapsulation on PTX Series Packet Transport Routers</i>• <i>Configuring CCC Encapsulation for Layer 2 VPNs</i>• <i>Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits</i>• <i>Configuring ATM for Subscriber Access</i>• <i>CoS on ATM IMA Pseudowire Interfaces Overview</i>• <i>Configuring Policing on an ATM IMA Pseudowire</i> |
|------------------------------|--|

encapsulation (Physical Interface)

Syntax	encapsulation (atm-ccc-cell-relay atm-pvc cisco-hdlc cisco-hdlc-ccc cisco-hdlc-tcc ethernet-bridge ethernet-ccc ethernet-over-atm ethernet-tcc ethernet-vpls ethernet-vpls-fr ether-vpls-over-atm-llc ethernet-vpls-ppp extended-frame-relay-ccc extended-frame-relay-ether-type-tcc extended-frame-relay-tcc extended-vlan-bridge extended-vlan-ccc extended-vlan-tcc extended-vlan-vpls flexible-ethernet-services flexible-frame-relay frame-relay frame-relay-ccc frame-relay-ether-type frame-relay-ether-type-tcc frame-relay-port-ccc frame-relay-tcc generic-services multilink-frame-relay-uni-nni ppp ppp-ccc ppp-tcc vlan-ccc vlan-vci-ccc vlan-vpls);
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces rlsq <i>number:number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers (flexible-ethernet-services , ethernet-ccc , and ethernet-tcc options only).
Description	Specify the physical link-layer encapsulation type. Not all encapsulation types are supported on the switches. See the switch CLI.
Default	ppp —Use serial PPP encapsulation.
Options	<p>atm-ccc-cell-relay—Use ATM cell-relay encapsulation.</p> <p>atm-pvc—Use ATM PVC encapsulation.</p> <p>cisco-hdlc—Use Cisco-compatible High-Level Data Link Control (HDLC) framing.</p> <p>cisco-hdlc-ccc—Use Cisco-compatible HDLC framing on CCC circuits.</p> <p>cisco-hdlc-tcc—Use Cisco-compatible HDLC framing on TCC circuits for connecting different media.</p> <p>ethernet-bridge—Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.</p> <p>ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces that must accept packets carrying standard Tag Protocol ID (TPID) values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, CCC is not supported.</p> <p>ethernet-over-atm—For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces. As defined in RFC 2684, <i>Multiprotocol Encapsulation over ATM Adaptation Layer 5</i>, this encapsulation type allows ATM interfaces to connect to devices that support only bridge protocol data units (BPDUs). Junos OS does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination</p>

IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.

ethernet-tcc—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.

ethernet-vpls-fr—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 Frame Relay connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use the MAC address to forward the packet into a given VPLS instance.

ethernet-vpls-ppp—Use in a VPLS setup when a CE device is connected to a PE device over a time division multiplexing (TDM) link. This encapsulation type enables the PE device to terminate the outer Layer 2 PPP connection, use the 802.1p bits inside the inner Ethernet header to classify the packets, look at the MAC address from the Ethernet header, and use it to forward the packet into a given VPLS instance.

ether-vpls-over-atm-llc—For ATM intelligent queuing (IQ) interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

extended-frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC.

extended-frame-relay-ether-type-tcc—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation type is used for circuits with different media on either side of the connection.

extended-frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.

extended-vlan-bridge—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

extended-vlan-ccc—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.

extended-vlan-tcc—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

extended-vlan-vpls—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

flexible-ethernet-services—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

flexible-frame-relay—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.

frame-relay—Use Frame Relay encapsulation.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with the Cisco Frame Relay.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect different media.

frame-relay-port-ccc—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect different media.

generic-services—Use generic services encapsulation for services with a hierarchical scheduler.

multilink-frame-relay-uni-nni—Use MLFR UNI NNI encapsulation. This encapsulation is used on link services, voice services interfaces functioning as FRF.16 bundles, and their constituent T1 or E1 interfaces, and is supported on LSQ and redundant LSQ interfaces.

ppp—Use serial PPP encapsulation.

ppp-ccc—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

ppp-tcc—Use serial PPP encapsulation on TCC circuits for connecting different media. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-ccc—Use Ethernet VLAN encapsulation on CCC circuits.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.

vlan-vpls—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only. On M Series routers, except the M320 router, the 4-port Fast Ethernet TX PIC and the 1-port, 2-port, and 4-port, 4-slot Gigabit Ethernet PICs can use the Ethernet VPLS encapsulation type.



NOTE:

- Label-switched interfaces (LSIs) do not support VLAN VPLS encapsulation. Therefore, you can only use VLAN VPLS encapsulation on a PE-router-to-CE-router interface and not a core-facing interface.
 - Starting with Junos OS release 13.3, a commit error occurs when you configure **vlan-vpls** encapsulation on a physical interface and configure **family inet** on one of the logical units. Previously, it was possible to commit this invalid configuration.
-

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
---------------------------------	---

**Related
Documentation**

- *Configuring Interface Encapsulation on Physical Interfaces*
- *Configuring CCC Encapsulation for Layer 2 VPNs*
- [Configuring Layer 2 Switching Cross-Connects Using CCC on page 647](#)
- *Configuring TCC Encapsulation for Layer 2 VPNs and Layer 2 Circuits*
- *Configuring ATM Interface Encapsulation*
- *Configuring ATM-to-Ethernet Interworking*
- *Configuring VLAN Encapsulation*
- *Configuring Extended VLAN Encapsulation*
- *Configuring Encapsulation for Layer 2 Wholesale VLAN Interfaces*
- *Configuring Interfaces for Layer 2 Circuits*
- *Configuring Interface Encapsulation on PTX Series Packet Transport Routers*
- [Configuring MPLS LSP Tunnel Cross-Connects Using CCC on page 655](#)
- [Configuring TCC on page 659](#)
- *Configuring VPLS Interface Encapsulation*
- *Configuring Interfaces for VPLS Routing*
- *Defining the Encapsulation for Switching Cross-Connects*

interface-switch

Syntax	<code>interface-switch <i>connection-name</i> { interface <i>interface-name.unit-number</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols connections], [edit protocols connections]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Configure Layer 2 switching cross-connects. The cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first.</p> <p>For Layer 2 switching cross-connects to work, you must also configure MPLS.</p>
Options	<p><i>connection-name</i>—Connection name (up to 128 characters in Junos 12.3 and later).</p> <p><i>interface interface-name.unit-number</i>—Interface name. Include the logical portion of the name, which corresponds to the logical unit number.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration</p>
Related Documentation	<ul style="list-style-type: none">• Configuring the CCC Connection for Layer 2 Switching Cross-Connects on page 652• <i>Defining the Connection for Switching Cross-Connects</i>• <i>MPLS Applications Feature Guide for Routing Devices</i>

lsp-switch

Syntax	<code>lsp-switch <i>connection-name</i> { transmit-lsp <i>label-switched-path</i>; receive-lsp <i>label-switched-path</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols connections], [edit protocols connections]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure Layer 2 switching cross-connects.
Options	<i>connection-name</i> —Connection name. <i>receive-lsp label-switched-path</i> —Name of the LSP from the connection's source. <i>transmit-lsp label-switched-path</i> —Name of the LSP to the connection's destination.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Connection for Layer 2 Switching TCCs on page 663

output-interface (CCC)

Syntax	<code>output-interface [<i>interface-name 1 interface-name n</i>];</code>
Hierarchy Level	[edit protocols connections p2mp-transmit-switch <i>p2mp-transmit-switch-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	Specify one or more output interfaces to switch traffic on an incoming CCC interface to one or more outgoing CCC interfaces.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring CCC Switching for Point-to-Multipoint LSPs on page 665

p2mp-receive-switch

Syntax	<pre>p2mp-receive-switch <i>point-to-multipoint-switch-name</i> { output-interface [<i>interface-name.unit-number</i>]; receive-p2mp-lsp <i>receiving-point-to-multipoint-lsp</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols connections], [edit protocols connections]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the CCC switch for a point-to-multipoint LSP on the egress PE router.
Options	<p><i>point-to-multipoint-switch-name</i>—Point-to-multipoint CCC receive switch name.</p> <p><i>output-interface interface-name.unit-number</i>—Name of the egress interfaces for the point-to-multipoint LSP traffic. You can configure multiple output interfaces.</p> <p><i>receive-p2mp-lsp receiving-point-to-multipoint-lsp</i>—Name of the point-to-multipoint LSP that is switched to the output interface.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Point-to-Multipoint LSP Switch on Egress PE Routers on page 667

p2mp-transmit-switch

Syntax	<code>p2mp-transmit-switch <i>point-to-multipoint-transmit-switch-name</i> { input-interface <i>interface-name.unit-number</i>; transmit-p2mp-lsp <i>transmitting-point-to-multipoint-lsp</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols connections], [edit protocols connections]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the CCC switch for the point-to-multipoint LSP on the ingress PE router.
Options	<p><i>point-to-multipoint-transmit-switch-name</i>—Point-to-multipoint CCC transmit switch name.</p> <p><i>input-interface interface-name.unit-number</i>—Specify the name of the interface carrying incoming traffic to be switched to the point-to-multipoint LSP.</p> <p><i>transmit-p2mp-lsp transmitting-point-to-multipoint-lsp</i>—Specify the name of the point-to-multipoint LSP carrying traffic to the CCC switch on the egress PE router.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring the Point-to-Multipoint LSP Switch on Ingress PE Routers on page 666

remote-interface-switch

Syntax	<pre>remote-interface-switch <i>connection-name</i> { interface <i>interface-name.unit-number</i>; transmit-lsp <i>label-switched-path</i>; receive-lsp <i>label-switched-path</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols connections], [edit protocols connections]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure MPLS LSP tunnel cross-connects.
Options	<p><i>connection-name</i>—Connection name.</p> <p><i>interface interface-name.unit-number</i>—Interface name. Include the logical portion of the name, which corresponds to the logical unit number.</p> <p><i>receive-lsp label-switched-path</i>—Name of the LSP from the connection's source.</p> <p><i>transmit-lsp label-switched-path</i>—Name of the LSP to the connection's destination.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MPLS LSP Tunnel Cross-Connects Using CCC on page 655

CHAPTER 27

GMPLS Configuration Statements

- [address \(Peer\) on page 1102](#)
- [control-channel \(Protocols Link Management Peer\) on page 1102](#)
- [dead-interval on page 1103](#)
- [disable \(GMPLS\) on page 1104](#)
- [disable \(OSPF\) on page 1105](#)
- [hello-dead-interval on page 1106](#)
- [hello-interval \(LMP\) on page 1107](#)
- [hello-interval \(Protocols OSPF\) on page 1108](#)
- [interface \(Protocols Link Management\) on page 1109](#)
- [label-switched-path \(Protocols Link Management\) on page 1109](#)
- [link-management on page 1110](#)
- [lmp-control-channel on page 1111](#)
- [lmp-protocol on page 1111](#)
- [local-address \(Protocols Link Management\) on page 1112](#)
- [passive \(Protocols Link Management\) on page 1112](#)
- [peer \(Protocols LMP\) on page 1113](#)
- [peer-interface \(Protocols OSPF\) on page 1114](#)
- [remote-address \(for LMP Control Channel\) on page 1115](#)
- [remote-address \(for LMP Traffic Engineering\) on page 1115](#)
- [remote-id on page 1116](#)
- [retransmission-interval on page 1116](#)
- [retransmit-interval \(OSPF\) on page 1117](#)
- [retry-limit \(Protocols Link Management\) on page 1118](#)
- [te-link on page 1119](#)
- [traceoptions \(Protocols Link Management\) on page 1120](#)
- [transit-delay \(OSPF\) on page 1122](#)
- [upstream-label on page 1123](#)

address (Peer)

Syntax	<code>address <i>ip-address</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i>],</code> <code>[edit protocols link-management peer <i>peer-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the ID of the peer.
Default	The loopback address is advertised.
Options	<i>ip-address</i> —IP address of the peer.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the ID for LMP Peers on page 681

control-channel (Protocols Link Management Peer)

Syntax	<code>control-channel <i>control-channel-interface</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i>],</code> <code>[edit protocols link-management peer <i>peer-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the control channel interface for the peer.
Options	<i>control-channel-interface</i> —Name of the control channel interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring LMP Peers on page 680

dead-interval

Syntax	<code>dead-interval seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Specify how long OSPF waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor.
Options	<p>seconds—Interval to wait.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: Four times the hello interval—40 seconds (broadcast and point-to-point networks); 120 seconds (nonbroadcast multiple access (NBMA) networks)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring OSPF Timers</i> • Configuring RSVP and OSPF for LMP Peer Interfaces on page 685

- [hello-interval on page 1108](#)

disable (GMPLS)

Syntax	disable;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i>], [edit protocols link-management te-link <i>te-link-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Disable a traffic engineering link.
Default	The configured object is enabled (operational) unless explicitly disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Disabling the Traffic Engineering Link for LMP Peers on page 685

disable (OSPF)

Syntax	disable;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> <i>peer-interface</i><i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instances</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3)],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) virtual-link],</p> <p>[edit protocols ospf area <i>area-id</i> <i>peer-interface</i> <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> virtual-link neighbor-id <i>router-id</i> transit-area <i>area-id</i>],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Disable OSPF, an OSPF interface, or an OSPF virtual link.

By default, control packets sent to the remote end of a virtual link must be forwarded using the default topology. In addition, the transit area path consists only of links that are in the default topology. You can disable a virtual link for a configured topology, but not for a default topology. Include the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id router-id transit-area *area-id* topology *name*]** hierarchy level.



NOTE: If you disable the virtual link by including the **disable** statement at the **[edit protocols ospf area *area-id* virtual-link neighbor-id router-id transit-area *area-id*]** hierarchy level, you disable the virtual link for all topologies, including the default topology. You cannot disable the virtual link only in the default topology.

Default	The configured object is enabled (operational) unless explicitly disabled.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding OSPF Configurations • Configuring RSVP and OSPF for LMP Peer Interfaces on page 685

hello-dead-interval

Syntax	<code>hello-dead-interval <i>milliseconds</i>;</code>
Hierarchy Level	<code>[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i> lmp-protocol],</code> <code>[edit protocols link-management peer <i>peer-name</i> lmp-protocol]</code>
Release Information	Statement introduced in Junos OS Release 8.0.
Description	Specify how long the Link Management Protocol (LMP) waits before declaring the control channel to be dead. This is an interval during which the router receives no LMP hello packets from the neighbor on a control that is active or up.
Options	<p><i>milliseconds</i>—Interval to wait before declaring the control channel to be dead.</p> <p>Range: 500 through 300,000</p> <p>Default: 500 milliseconds (three times the hello interval)</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Hello Message Intervals for LMP Control Channels on page 683 • hello-interval (LMP) on page 1107

hello-interval (LMP)

Syntax	hello-interval <i>milliseconds</i> ;
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i> <i>lmp-protocol</i>], [edit protocols link-management peer <i>peer-name</i> <i>lmp-protocol</i>]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	Specify how often the router sends Link Management Protocol (LMP) hello packets.
Options	<i>milliseconds</i> —Length of time between hello packets. Range: 150 through 300,000 Default: 150 milliseconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Hello Message Intervals for LMP Control Channels on page 683• hello-dead-interval on page 1106

hello-interval (Protocols OSPF)

Syntax	hello-interval <i>seconds</i> ;
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Specify how often the routing device sends hello packets out the interface. The hello interval must be the same for all routing devices on a shared logical IP network.
Options	<p>seconds—Time between hello packets, in seconds.</p> <p>Range: 1 through 255 seconds</p> <p>Default: 10 seconds (broadcast and point-to-point networks); 30 seconds (nonbroadcast multiple access [NBMA] networks)</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring OSPF Timers</i> • Configuring RSVP and OSPF for LMP Peer Interfaces on page 685 • dead-interval on page 1103

interface (Protocols Link Management)

Syntax	<code>interface <i>interface-name</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i>], [edit protocols link-management te-link <i>te-link-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the egress router interface.
Options	<i>interface-name</i> —Name of the interface to the egress router.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • LMP Configuration Overview on page 677

label-switched-path (Protocols Link Management)

Syntax	<code>label-switched-path <i>lsp-name</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i>], [edit protocols link-management te-link <i>te-link-name</i>]
Release Information	Statement introduced in Junos OS Release 7.4.
Description	Specify the label-switched path (LSP) to be used by the forwarding adjacency.
Options	<i>lsp-name</i> —Name of the LSP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Forwarding Adjacencies on page 725

link-management

Syntax	<pre> link-management { peer <i>peer-name</i> { address <i>ip-address</i>; control-channel <i>control-channel-interface</i>; lmp-control-channel <i>control-channel-interface</i> { remote-address <i>ip-address</i>; } lmp-protocol { hello-dead-interval <i>milliseconds</i>; hello-interval <i>milliseconds</i>; passive; retransmission-interval <i>milliseconds</i>; retry-limit <i>number</i>; } te-link <i>te-link-name</i>; } te-link <i>te-link-name</i> { disable; interface <i>interface-name</i> { disable; local-address <i>ip-address</i>; remote-address <i>ip-address</i>; remote-id <i>id-number</i>; } local-address <i>ip-address</i>; remote-address <i>ip-address</i>; remote-id <i>id-number</i>; } traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; } } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols], [edit protocols]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable Link Management Protocol (LMP) on the router.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • LMP Configuration Overview on page 677

lmp-control-channel

Syntax	<code>lmp-control-channel <i>control-channel-interface</i> { <i>remote-address</i> <i>ip-address</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i>], [edit protocols link-management peer <i>peer-name</i>]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	Specify the Link Management Protocol (LMP) control channel interface for the peer.
Options	<i>control-channel-interface</i> —Name of the control channel interface. The remaining statement is described separately in this chapter.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the LMP Control Channel Interface for the Peer on page 681

lmp-protocol

Syntax	<code>lmp-protocol { <i>hello-dead-interval</i> <i>milliseconds</i>; <i>hello-interval</i> <i>milliseconds</i>; passive; <i>retransmission-interval</i> <i>milliseconds</i>; <i>retry-limit</i> <i>number</i>; }</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i>], [edit protocols link-management peer <i>peer-name</i>]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	Configure attributes of Link Management Protocol (LMP) to establish and maintain the LMP control channel for the peer.
Options	The statements are described separately in this chapter.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring LMP Peers on page 680

local-address (Protocols Link Management)

Syntax	<code>local-address <i>ip-address</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i>], [edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i> interface <i>interface-name</i>], [edit protocols link-management te-link <i>te-link-name</i>], [edit protocols link-management te-link <i>te-link-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the local IP address associated with the traffic engineering link. If you configure the local IP address, you must also configure the remote-address statement.
Options	local-address —Local IP address of the traffic engineering link.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Local IP Address for Traffic Engineering Links on page 679• Configuring the Local IP Address for Forwarding Adjacencies on page 725• remote-address (for LMP Traffic Engineering) on page 1115

passive (Protocols Link Management)

Syntax	<code>passive;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i> lmp-protocol], [edit protocols link-management peer <i>peer-name</i> lmp-protocol]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	Specify that the router not configure the Link Management Protocol (LMP) control channels but wait for the remote peer to configure the LMP control channels.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Preventing the Local Peer from Initiating LMP Negotiation on page 684

peer (Protocols LMP)

Syntax	<pre> peer <i>peer-name</i> { address <i>ip-address</i>; control-channel <i>control-channel-interface</i>; lmp-control-channel <i>control-channel-interface</i>; lmp-protocol { hello-dead-interval <i>milliseconds</i>; hello-interval <i>milliseconds</i>; passive; retransmission-interval <i>milliseconds</i>; retry-limit <i>number</i>; } te-link <i>te-link-name</i>; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management], [edit protocols link-management]
Release Information	Statement introduced before Junos OS Release 7.4. lmp-protocol statement and substatements added in Junos OS Release 8.1.
Description	Configure a network peer.
Options	<p><i>peer-name</i>—Name of the network peer.</p> <p>The remaining statements are described separately in this chapter.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring LMP Peers on page 680

peer-interface (Protocols OSPF)

Syntax	<pre>peer-interface <i>interface-name</i> { disable; dead-interval <i>seconds</i>; hello-interval <i>seconds</i>; retransmit-interval <i>seconds</i>; transit-delay <i>seconds</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i>], [edit protocols ospf area <i>area-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure a peer interface.
Options	<p><i>interface-name</i>—Name of the peer interface. To configure all interfaces, you can specify <i>all</i>.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring OSPFv2 Peer interfaces for GMPLS• Configuring RSVP and OSPF for LMP Peer Interfaces on page 685• Advertising Forwarding Adjacencies Using OSPF on page 727

remote-address (for LMP Control Channel)

Syntax	<code>remote-address <i>ip-address</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i> lmp-control-channel <i>control-channel-interface</i>], [edit protocols link-management peer <i>peer-name</i> lmp-control-channel <i>control-channel-interface</i>]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	Specify the remote IP address for the Link Management Protocol (LMP) control channel interface.
Options	<i>ip-address</i> —Remote IP address mapped to the LMP control channel interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Remote IP Address for LMP Control Channels on page 682

remote-address (for LMP Traffic Engineering)

Syntax	<code>remote-address <i>ip-address</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i>], [edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i> interface <i>interface-name</i>], [edit protocols link-management te-link <i>te-link-name</i>], [edit protocols link-management te-link <i>te-link-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the remote IP address for the traffic engineering link. If you configure the remote IP address, you must also configure the local-address statement.
Options	<i>ip-address</i> —Remote IP address mapped to the traffic engineering link.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Remote IP Address for Traffic Engineering Links on page 679 • Configuring the Remote IP Address for Forwarding Adjacencies on page 726 • local-address (Protocols Link Management) on page 1112

remote-id

Syntax	<code>remote-id <i>id-number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i>], [edit logical-systems <i>logical-system-name</i> protocols link-management te-link <i>te-link-name</i> interface <i>interface-name</i>], [edit protocols link-management te-link <i>te-link-name</i>], [edit protocols link-management te-link <i>te-link-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the ID assigned to a traffic engineering link or an interface (resource) on the peer node.
Options	<i>id-number</i> —ID number for the remote device.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Remote ID for Traffic Engineering Links on page 680

retransmission-interval

Syntax	<code>retransmission-interval <i>milliseconds</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i> lmp-protocol], [edit protocols link-management peer <i>peer-name</i> lmp-protocol]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	Specify how often Link Management Protocol (LMP) sends Config and LinkSummary messages on the LMP control channel.
Options	<i>milliseconds</i> —Length of time between Config messages. Range: 500 through 300,000 Default: 500 milliseconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• retry-limit (Protocols Link Management) on page 1118• Controlling Message Exchange for LMP Control Channels on page 684

retransmit-interval (OSPF)

Syntax	<code>retransmit-interval seconds;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	Specify how long the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements (LSAs) to an interface's neighbors.
Options	<p>seconds—Interval to wait.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 5 seconds</p>



NOTE: You must configure LSA retransmit intervals to be equal to or greater than 3 seconds to avoid triggering a retransmit trap, because Junos OS delays LSA acknowledgments by up to 2 seconds.

Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring OSPF Timers• Configuring RSVP and OSPF for LMP Peer Interfaces on page 685

retry-limit (Protocols Link Management)

Syntax	<code>retry-limit <i>number</i>;</code>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management peer <i>peer-name</i> lmp-protocol], [edit protocols link-management peer <i>peer-name</i> lmp-protocol]
Release Information	Statement introduced in Junos OS Release 8.1.
Description	Specify how many times the Link Management Protocol (LMP) sends Config and LinkSummary messages on the LMP control channel without receiving an appropriate acknowledgment before it logs a message and restarts the LMP control channel configuration process.
Options	<i>number</i> —Maximum number of times messages are sent without receiving an acknowledgment. Range: 3 through 1000 Default: 3
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• retransmission-interval on page 1116• Controlling Message Exchange for LMP Control Channels on page 684

te-link

Syntax	<pre>te-link <i>te-link-name</i> { disable; ethernet-vlan; interface <i>interface-name</i> { disable; local-address <i>ip-address</i>; remote-address <i>ip-address</i>; remote-id <i>id-number</i>; } local-address <i>ip-address</i>; remote-address <i>ip-address</i>; remote-id <i>id-number</i>; }</pre>
Hierarchy Level	[edit protocols link-management], [edit protocols link-management peer <i>peer-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. ethernet-vlan option introduced in Junos OS Release 14.2.
Description	Represent a collection of physical ports or time slots. Assign a traffic engineering link to the specified network peer.
Options	<p><i>te-link-name</i>—Name of the collection of physical ports or the name of the time slots.</p> <p>disable—Disable the traffic engineering link or an interface to a traffic engineering link.</p> <p>The remaining statements are explained separately..</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring LMP Traffic Engineering Links on page 678

traceoptions (Protocols Link Management)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> <flag-modifier> <disable>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols link-management], [edit protocols link-management]
Release Information	Statement introduced before Junos OS Release 7.4. Support for hello-packets , packets , and state flags added in Junos OS Release 8.1.
Description	Trace options for the LMP protocol.
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000</p> <p>Default: 2 files</p> <p>If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.</p> <p>flag—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.</p> <ul style="list-style-type: none">• all—Trace all available operations• hello-packets—Trace hello packets on any LMP control channel• init—Output from the initialization messages• packets—Trace all packets other than hello packets on any LMP control channel• parse—Operation of the parser• process—Operation of the general configuration• route-socket—Operation of route socket events• routing—Operation of the routing protocols• server—Server processing operations• show—show command servicing operations

- **state**—Trace state transitions of the LMP control channels and traffic engineering links

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-world-readable—(Optional) Prevent all users from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Enable log file access for all users.

Required Privilege Level	routing and trace—To view this statement in the configuration. routing-control and trace-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Tracing LMP Traffic on page 687 • <i>Network Management Administration Guide for Routing Devices</i>

transit-delay (OSPF)

Syntax	<code>transit-delay <i>seconds</i>;</code>
Hierarchy Level	<p>[edit logical-systems <i>logical-system-name</i> protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols ospf area <i>area-id</i> peer-interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit protocols (ospf ospf3) area <i>area-id</i> virtual-link],</p> <p>[edit protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast)] area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf area <i>area-id</i> virtual-link],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ospf3 realm (ipv4-unicast ipv4-multicast ipv6-multicast) area <i>area-id</i> interface <i>interface-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2.</p> <p>Support for the realm statement introduced in Junos OS Release 9.2 for EX Series switches.</p>
Description	<p>Set the estimated time required to transmit a link-state update on the interface. When calculating this time, make sure to account for transmission and propagation delays.</p> <p>You should never have to modify the transit delay time.</p>
Options	<p>seconds—Estimated time, in seconds.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: 1 second</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring OSPF Timers</i> • Configuring RSVP and OSPF for LMP Peer Interfaces on page 685

upstream-label

Syntax	<pre>upstream-label { vlan-id <i>vlan-id</i>; }</pre>
Hierarchy Level	[edit protocols mpls label-switched-path <i>lsp-name</i> lsp-attributes]
Release Information	Statement introduced in Junos OS Release 14.2.
Description	Specify the upstream label for the bidirectional label-switched path (LSP).
Options	vlan-id <i>vlan-id</i> —VLAN ID to be used for the Generalized MPLS (GMPLS) VLAN LSP at the ingress provider edge (PE) to customer edge (CE) link.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MPLS LSPs for GMPLS on page 688

CHAPTER 28

PCEP Configuration Statements

- [\[edit protocols pcep\] Hierarchy Level](#) on page 1125
- [pcep](#) on page 1127
- [delegation-cleanup-timeout](#) on page 1128
- [delegation-priority](#) on page 1128
- [destination-ipv4-address](#) on page 1129
- [destination-port](#) on page 1129
- [lsp-external-controller](#) on page 1130
- [max-unknown-messages](#) on page 1130
- [message-rate-limit](#) on page 1131
- [pce](#) on page 1132
- [pce-group \(PCE\)](#) on page 1133
- [pce-group \(Protocols PCEP\)](#) on page 1134
- [pce-type](#) on page 1135
- [querier \(performance-monitoring\)](#) on page 1136
- [traceoptions \(PCE\)](#) on page 1137
- [traceoptions \(Protocols PCEP\)](#) on page 1139
- [update-rate-limit](#) on page 1140

[\[edit protocols pcep\] Hierarchy Level](#)

```
pcep {  
  message-rate-limit messages-per-minute;  
  pce pce-id {  
    delegation-cleanup-timeout seconds;  
    delegation-priority priority-number;  
    destination-ipv4-address ipv4-address;  
    destination-port port-number;  
    max-unknown-messages messages-per-minute;  
    pce-group pce-group-name;  
    pce-type {  
      active stateful;  
    }  
    traceoptions {
```

```
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag (all | pcep);
        no-remote-trace;
    }
}
pce-group pce-group-id {
    delegation-cleanup-timeout seconds;
    max-unknown-messages messages-per-minute;
    pce-type {
        active stateful;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag (all | pcep);
        no-remote-trace;
    }
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
update-rate-limit updates-per-minute;
}
```

pcep

Syntax	<pre> pcep { message-rate-limit <i>messages-per-minute</i>; pce <i>pce-id</i> { delegation-cleanup-timeout <i>seconds</i>; delegation-priority <i>priority-number</i>; destination-ipv4-address <i>ipv4-address</i>; destination-port <i>port-number</i>; max-unknown-messages <i>messages-per-minute</i>; pce-group <i>pce-group-name</i>; pce-type { active stateful; } traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag (all pcep); no-remote-trace; } } pce-group <i>pce-group-id</i> { delegation-cleanup-timeout <i>seconds</i>; max-unknown-messages <i>messages-per-minute</i>; pce-type { active stateful; } traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag (all pcep); no-remote-trace; } } traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; no-remote-trace; } update-rate-limit <i>updates-per-minute</i>; } </pre>
Hierarchy Level	[edit protocols]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	<p>Configure the Path Computation Client (PCC) parameters.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Path Computation Element Protocol for MPLS RSVP-TE on page 733

delegation-cleanup-timeout

Syntax	<code>delegation-cleanup-timeout <i>seconds</i>;</code>
Hierarchy Level	<code>[edit protocols pcep pce <i>pce-id</i>]</code>
Release Information	Statement introduced in Junos OS Release 12.3. Support for PTX Series added in Junos OS Release 14.2.
Description	Specify the amount of time (in seconds) that a Path Computation Client (PCC) must wait before returning control of all LSPs to the routing protocol process after a PCEP session with the main active stateful Path Computation Element (PCE) is disconnected.
Options	<i>seconds</i> —Time (in seconds) that a PCC must wait before returning control of all LSPs to the routing protocol process after a PCEP session with the main active stateful PCE is disconnected. Range: 0 through 600 seconds Default: 30 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• pce on page 1132

delegation-priority

Syntax	<code>delegation-priority <i>priority-number</i>;</code>
Hierarchy Level	<code>[edit protocols pcep pce <i>pce-id</i>]</code>
Release Information	Statement introduced in Junos OS Release 12.3. Support for PTX Series added in Junos OS Release 14.2.
Description	Specify the priority number of the active stateful Path Computation Element (PCE). This value is used by the Path Computation Client (PCC) to elect a PCE to delegate LSPs. No two PCEs can have the same delegation-priority value. The PCC elects the PCE with a lower priority as the main active stateful PCE to delegate LSPs.
Options	<i>priority-number</i> —Priority number of the active stateful PCE. Range: 1 through 65535 Default: 0 (no priority is set)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• pce on page 1132

destination-ipv4-address

Syntax	<code>destination-ipv4-address <i>ipv4-address</i>;</code>
Hierarchy Level	[edit protocols pcep pce <i>pce-id</i>]
Release Information	Statement introduced in Junos OS Release 12.3. Support for PTX Series added in Junos OS Release 14.2.
Description	Specify the IPv4 address of the Path Computation Element (PCE) to which the Path Computation Client (PCC) should connect.
Options	<i>ipv4-address</i> —IPv4 address of the PCE.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• pce on page 1132

destination-port

Syntax	<code>destination-port <i>port-number</i>;</code>
Hierarchy Level	[edit protocols pcep pce <i>pce-id</i>]
Release Information	Statement introduced in Junos OS Release 12.3. Support for PTX Series added in Junos OS Release 14.2.
Description	Specify the TCP port number of the Path Computation Element (PCE) to which the Path Computation Client (PCC) should connect.
Options	<i>port-number</i> —Destination TCP port number. Range: 1 through 65535 Default: 4189
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• pce on page 1132

`lsp-external-controller`

Syntax	<code>lsp-external-controller <i>lsp-external-controller</i>;</code>
Hierarchy Level	[edit protocols mpls], [edit protocols mpls label-switched-path <i>lsp-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3.
Description	Configure an external path computing entity.
Options	<i>lsp-external-controller</i> —Name of the external path computing entity. Values: <i>pccd</i>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• pcep on page 1127

`max-unknown-messages`

Syntax	<code>max-unknown-messages <i>messages-per-minute</i>;</code>
Hierarchy Level	[edit protocols pcep pce <i>pce-id</i>]
Release Information	Statement introduced in Junos OS Release 12.3. Support for PTX Series added in Junos OS Release 14.2.
Description	Specify the number of unknown messages per minute that the Path Computation Client (PCC) can receive at maximum after which the PCEP session is closed.
Options	<i>messages-per-minute</i> —Number of unknown messages per minute that the PCC can receive at maximum after which the PCEP session is closed. Recommended value is 5. If the number of unknown messages received by the PCC is greater than or equal to the maximum number, the PCEP session is closed. Range: 1 through 16384 Default: 0 (disabled or no limit)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• pce on page 1132

message-rate-limit

Syntax	<code>message-rate-limit <i>messages-per-minute</i>;</code>
Hierarchy Level	[edit protocols pcep]
Release Information	Statement introduced in Junos OS Release 12.3. Support for PTX Series added in Junos OS Release 14.2.
Description	Specify the number of messages per minute that the Path Computation Client (PCC) can receive at maximum.
Options	<i>messages-per-minute</i> —Number of messages per minute that the PCC can receive at maximum. Range: 1 through 16384 Default: 0 (disabled or no limit)
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• pcep on page 1127

pce

Syntax `pce pce-id {
 delegation-cleanup-timeout seconds;
 delegation-priority priority-number;
 destination-ipv4-address ipv4-address;
 destination-port port-number;
 max-unknown-messages messages-per-minute;
 pce-group pce-group-name;
 pce-type {
 active stateful;
 }
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag (all | pcep);
 no-remote-trace;
 }
 }
 }`

Hierarchy Level [edit protocols pcep]

Release Information Statement introduced in Junos OS Release 12.3.
 Support for PTX Series added in Junos OS Release 14.2.

Description Configure per path computation element (PCE) parameters.

Options *pce-id*—IP address of the PCE.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation • [pcep on page 1127](#)

pce-group (PCE)

Syntax	<code>pce-group <i>pce-group-name</i>;</code>
Hierarchy Level	[edit protocols pcep pce <i>pce-id</i>]
Release Information	Statement introduced in Junos OS Release 12.3. Support for PTX Series added in Junos OS Release 14.2.
Description	Specify the Path Computation Element (PCE) group to which the configured PCE belongs.
Options	<i>pce-group-name</i> —Name of the PCE group.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• pce on page 1132

pce-group (Protocols PCEP)

Syntax `pce-group pce-group-id {
 delegation-cleanup-timeout seconds;
 max-unknown-messages messages-per-minute;
 pce-type {
 active stateful;
 }
 traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag (all | pcep);
 no-remote-trace;
 }
}`

Hierarchy Level [edit protocols pcep]

Release Information Statement introduced in Junos OS Release 12.3.
Support for PTX Series added in Junos OS Release 14.2.

Description Configure the Path Computation Element (PCE) group parameters. A maximum of 10 PCE groups can be configured at any given point in time.
The remaining statements are explained separately.



NOTE: A PCE group can include PCEs that are either only stateful or only active stateful. A combination of stateful PCEs and active stateful PCEs in one PCE group is not supported.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation • [pcep on page 1127](#)

pce-type

Syntax	<pre>pce-type { active stateful; }</pre>
Hierarchy Level	[edit protocols pcep pce <i>pce-id</i>]
Release Information	Statement introduced in Junos OS Release 12.3. Support for PTX Series added in Junos OS Release 14.2.
Description	<p>Configure the path computation element (PCE) type:</p> <ul style="list-style-type: none">• active—Uses LSP state information learned from PCCs to optimize path computations, and actively updates LSP parameters in those PCCs that delegated control over their LSPs to the PCE.• stateful—Uses LSP state information learned from PCCs to optimize path computations, but does not actively update the LSP state. A PCC maintains synchronization with the PCE.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• pce on page 1132

querier (performance-monitoring)

```
Syntax  querier {
        delay {
            traffic-class tc-value {
                average-sample-size sample size;
                padding-size size;
                query-interval milliseconds;
                rtt-delay-threshold rtt threshold value;
                twcd-delay-threshold twcd threshold value;
            }
        }
        loss {
            traffic-class tc-value {
                average-sample-size sample size;
                loss-threshold loss threshold value;
                loss-threshold-window number of samples for loss threshold;
                measurement-quantity bytes|packets;
                query-interval milliseconds;
            }
        }
        loss-delay {
            traffic-class tc-value {
                average-sample-size sample size;
                loss-threshold loss threshold value;
                loss-threshold-window number of samples for loss threshold;
                measurement-quantity bytes|packets;
                padding-size size;
                query-interval milliseconds;
                rtt-delay-threshold rtt threshold value;
                twcd-delay-threshold twcd threshold value;
            }
        }
    }
```

Hierarchy Level [edit protocols mpls [oam performance-monitoring](#)],
 [edit protocols mpls [label-switched-path](#) *lsp-name* [oam performance-monitoring](#)],
 [edit protocols mpls [label-switched-path](#) *lsp-name* [primary path-name oam performance-monitoring](#)],
 [edit protocols mpls [label-switched-path](#) *lsp-name* [secondary path-name oam performance-monitoring](#)]

Release Information Statement introduced in Junos OS Release 15.1.

Description Configure querier options.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Documentation •

traceoptions (PCE)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag (all pcep); no-remote-trace; } </pre>
Hierarchy Level	[edit protocols pcep pce <i>pce-id</i>]
Description	Configure the Path Computation Element Protocol (PCEP) tracing options.
Options	<p><i>filename</i>—Name of the file to receive the output of the tracing operation. All files are placed in the directory <i>/var/log</i>.</p> <p><i>files number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files. If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.</p> <p><i>flag</i>—Area of path computation client process (pccd) to enable debugging output.</p> <ul style="list-style-type: none"> all—Trace all areas of PCD code. pcep—Trace Path Computation Element Protocol (PCEP) operations. <p><i>no-remote-trace</i>—(Optional) Disable remote tracing options.</p> <p><i>no-world-readable</i>—(Optional) Allow only certain users to read the log file.</p> <p><i>size size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches this size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Syntax: <i>xk</i> to specify KB, <i>xm</i> to specify MB, or <i>xg</i> to specify GB.</p> <p>Range: 10 KB through the maximum file size supported on your system.</p> <p>Default: 1 MB. If you specify a maximum file size, you must also include the files statement to specify the maximum number of files.</p> <p><i>world-readable</i>—(Optional) Allow any user to read the log file.</p>
Required Privilege	routing and trace—To view this statement in the configuration.
Level	routing-control and trace-control—To add this statement to the configuration.

Related • [pce on page 1132](#)
Documentation

traceoptions (Protocols PCEP)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; no-remote-trace; } </pre>
Hierarchy Level	[edit protocols pcep]
Description	Configure the Path Computation Element Protocol (PCEP) tracing options.
Options	<p><i>filename</i>—Name of the file to receive the output of the tracing operation. All files are placed in the directory <i>/var/log</i>.</p> <p><i>files number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Range: 2 through 1000 files</p> <p>Default: 2 files. If you specify a maximum number of files, you must also include the size statement to specify the maximum file size.</p> <p><i>flag</i>—Area of path computation client process (pccd) to enable debugging output.</p> <p>PCEP Tracing Flags</p> <ul style="list-style-type: none"> • <i>all</i>—Trace all areas of PCCD code • <i>pccd-config</i>—All configuration parsing operations • <i>pccd-core</i>—PCCD core operations • <i>pccd-functions</i>—PCCD function entries and outs • <i>pccd-main</i>—PCCD main module • <i>pccd-rpd</i>—PCCD communication with RPD • <i>pccd-ui</i>—PCCD user interface handling <p><i>no-remote-trace</i>—(Optional) Disable remote tracing options.</p> <p><i>no-world-readable</i>—(Optional) Allow only certain users to read the log file.</p> <p><i>size size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches this size, it is renamed trace-file.0. When the trace-file again reaches this size, trace-file.0 is renamed trace-file.1 and trace-file is renamed trace-file.0. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p>Syntax: <i>xk</i> to specify KB, <i>xm</i> to specify MB, or <i>xg</i> to specify GB.</p> <p>Range: 10 KB through the maximum file size supported on your system.</p>

Default: 1 MB. If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- [pcep on page 1127](#)

update-rate-limit

Syntax `update-rate-limit updates-per-minute;`

Hierarchy Level [edit protocols pcep]

Release Information Statement introduced in Junos OS Release 12.3.
Support for PTX Series added in Junos OS Release 14.2.

Description Specify the number of updates per minute that the Path Computation Client (PCC) can receive at maximum. Updates above this limit are ignored by the PCC.

Options ***updates-per-minute***—Number of updates per minute that the PCC can receive at maximum.
Range: 1 through 16384
Default: 0 (disabled or no limit)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [pcep on page 1127](#)

PART 10

Operational Commands

- [MPLS Operational Commands on page 1143](#)
- [RSVP Operational Commands on page 1259](#)
- [LDP Operational Commands on page 1303](#)
- [CCC and TCC Operational Commands on page 1355](#)
- [PCEP Operational Commands on page 1375](#)


CHAPTER 29

MPLS Operational Commands

- `clear mpls lsp`
- `clear mpls container-lsp`
- `clear performance-monitoring mpls lsp`
- `monitor mpls delay rsvp`
- `monitor mpls loss rsvp`
- `monitor mpls loss-delay rsvp`
- `ping mpls bgp`
- `ping mpls lsp-end-point`
- `request mpls container-lsp`
- `request mpls lsp adjust-autobandwidth`
- `show connections`
- `show link-management`
- `show link-management peer`
- `show link-management routing`
- `show link-management statistics`
- `show link-management te-link`
- `show mpls admin-groups`
- `show mpls call-admission-control`
- `show mpls container-lsp`
- `show mpls context-identifier`
- `show mpls cspf`
- `show mpls diffserv-te`
- `show mpls egress-protection`
- `show mpls interface`
- `show mpls label usage`
- `show mpls lsp`
- `show mpls lsp autobandwidth`
- `show mpls path`

- `show mpls srlg`
- `show mpls static-lsp`
- `show performance-monitoring mpls lsp`
- `show ted database`
- `show ted link`
- `show ted protocol`
- `traceroute mpls bgp`

clear mpls lsp

List of Syntax	Syntax on page 1145 Syntax (EX and QFX Series Switches) on page 1145
Syntax	<pre>clear mpls lsp <autobandwidth> <logical-system (all <i>logical-system-name</i>)> <name <i>name</i>> <optimize optimize-aggressive> <path <i>regular-expression</i>> <statistics></pre>
Syntax (EX and QFX Series Switches)	<pre>clear mpls lsp <autobandwidth> <name <i>name</i>> <optimize optimize-aggressive> <path <i>regular-expression</i>> <statistics></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p>
Description	Release the routes and states associated with MPLS label-switched paths (LSPs), and start new LSPs.
<div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p>CAUTION: This command disconnects existing Resource Reservation Protocol (RSVP) sessions on the ingress routing device. If there is a time lag between the old path being torn down and the new path being set up, this command might impact traffic traveling along the LSPs.</p> </div> </div>	
Options	<p>none—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.</p> <p>autobandwidth—(Optional) Clear LSP autobandwidth counters.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>name <i>name</i>—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the <i>Junos Network Interfaces Configuration Guide</i>.</p> <p>optimize optimize-aggressive—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.</p>

path *regular-expression*—(Optional) Clear the specific LSP path matching the specified regular expression.

statistics—(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (**name** and **path** options) on transit routers.

Required Privilege Level

clear

Related Documentation

- [show mpls lsp on page 1209](#)
- [show rsvp session on page 1281](#)

List of Sample Output [clear mpls lsp on page 1146](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[clear mpls lsp](#)

```
user@host> clear mpls lsp
```


clear mpls container-lsp

Syntax	<pre>clear mpls container-lsp <autobandwidth> <history> <logical-system (all <i>logical-system-name</i>)> <member> <name <i>name</i>> <optimize optimize-aggressive> <statistics></pre>
Release Information	Command introduced in Junos OS Release 14.2.
Description	Release the routes and states associated with MPLS container label-switched paths (LSPs), and start new LSPs.
Options	<p>none—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.</p> <p>autobandwidth—(Optional) Clear LSP autobandwidth counters.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>name <i>name</i>—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the <i>Junos Network Interfaces Configuration Guide</i>.</p> <p>optimize optimize-aggressive—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.</p> <p>statistics—(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (name and path options) on transit routers.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show mpls container-lsp on page 1190 • request mpls container-lsp on page 1167
List of Sample Output	clear mpls container-lsp on page 1148 clear mpls container-lsp name on page 1148 clear mpls container-lsp statistics on page 1148
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear mpls container-lsp

```
user@host> clear mpls container-lsp
```

clear mpls container-lsp name

```
user@host> clear mpls container-lsp name name
```

clear mpls container-lsp statistics

```
user@host> clear mpls container-lsp statistics
```

clear performance-monitoring mpls lsp

Syntax	clear performance-monitoring mpls lsp <name <i>lsp-name</i> >
Release Information	Command introduced in Junos OS Release 15.1.
Description	Restart the performance monitoring statistics.
Options	<p>none—Reset and restart all performance monitoring for all LSPs.</p> <p>name <i>lsp-name</i>—(Optional) Reset and restart performance monitoring for the specified LSP.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • performance-monitoring (Protocols MPLS) on page 922 • show performance-monitoring mpls lsp on page 1236
List of Sample Output	clear performance-monitoring mpls lsp on page 1149
Output Fields	When you enter this command, performance monitoring is restarted.

Sample Output

clear performance-monitoring mpls lsp

```
user@host> clear performance-monitoring mpls lsp
```

monitor mpls delay rsvp

Syntax	<code>monitor mpls delay rsvp <i>lsp-name</i></code> <code><detail></code> <code><count <i>count</i>></code> <code><interval <i>seconds</i>></code> <code><padding-size <i>padding-size</i>></code> <code><traffic-class <i>traffic-class</i>></code>
Release Information	Command introduced in Junos OS Release 14.2.
Description	Perform an on-demand delay measurement and display the measured values for associated bidirectional MPLS ultimate hop popping (UHP) point-to-point label-switched paths (LSPs).
Options	<p><i>lsp-name</i>—Name of the associated bidirectional MPLS UHP LSP for which the delay measurement is performed.</p> <p>detail—(Optional) Display detailed output of the LSP delay measurement.</p> <p>count <i>count</i>—(Optional) Specify the number of delay measurements to be carried out for the MPLS UHP LSP. For LSP delay measurement, the number of queries sent is the specified count number plus one additional query, because the LSP delay is measured using successive messages. Default: 10 Range: 1 through 1000000</p> <p>interval <i>seconds</i>—(Optional) Specify in seconds the interval between two successive query messages. Range: 1 through 255 seconds</p> <p>padding-size <i>padding-size</i>—(Optional) Specify the length of padding TLV to be included in the query message. Range: 0 through 1500</p> <p>traffic-class <i>traffic-class</i>—(Optional) Specify the traffic class for the LSP delay measurement. When the traffic-class value is not specified, the default traffic-class code-point of 111 is used. Range: 0 though 7</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• monitor mpls loss rsvp on page 1154• monitor mpls loss-delay rsvp on page 1159• Example: Configuring On-Demand Loss and Delay Measurement on page 439
List of Sample Output	monitor mpls lsp delay rsvp count on page 1152

[monitor mpls lsp delay rsvp count detail on page 1152](#)

Output Fields [Table 20 on page 1151](#) describes the output fields for the **monitor mpls delay rsvp** command. Output fields are listed in the approximate order in which they appear.

Table 20: monitor mpls delay rsvp Output Fields

Field Name	Field Description	Level of Output
Current two-way channel delay	Sum of packet delays, excluding the processing time of the remote provider edge (PE) router.	All Levels
Current round-trip-time	Total time taken for completing round-trip of packet.	All Levels
Best two-way channel delay	Best available two-way channel delay count.	All Levels
Worst two-way channel delay	Worst available two-way channel delay count.	All Levels
Average two-way channel delay	Average of the available two-way channel delay counts.	All Levels
Best round-trip-time	Best available round-trip-time count.	All Levels
Worst round-trip-time	Worst available round-trip-time count.	All Levels
Average round-trip-time	Average of the available round-trip-time counts.	All Levels
Average forward delay variation	Average of the variation in forward delay.	All Levels
Average reverse delay variation	Average of the variation in reverse delay.	All Levels
DM queries sent	Number of queries sent for delay measurement.	All Levels
DM responses received	Number of responses received for delay measurement queries.	All Levels
DM queries timedout	Number of timed out queries sent for delay measurement.	All Levels
DM responses dropped due to errors	Number of loss measurement responses dropped due to errors.	All Levels
Response code	Status of the messages used for delay measurement. Response code can be one of the following: <ul style="list-style-type: none"> • Success—Successful response code. • Failed—Failed response code. 	detail
Querier transmit timestamp	Timestamp on the query message when the message is sent out the ingress PE router (querier). This is done in the hardware before packet is sent out of an interface.	detail

Table 20: monitor mpls delay rsvp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Responder receive timestamp	Timestamp on the response message when the message is received by the egress PE router (responder). This is done in the hardware before packet is received by an interface.	detail
Responder transmit timestamp	Timestamp on the query message when the message is sent out the egress PE router (responder). This is done in the hardware before packet is sent out of an interface.	detail
Querier receive timestamp	Timestamp on the response message when the message is received by the ingress PE router (querier). This is done in the hardware before packet is received by an interface.	detail

Sample Output

monitor mpls lsp delay rsvp count

```

user@host> monitor mpls lsp delay rsvp LSP-A count 2

(1)
Current two-way channel delay      : 44 usecs
Current round-trip-time            : 3243 usecs
(2)
Current two-way channel delay      : 45 usecs
Current round-trip-time            : 1752 usecs

Best two-way channel delay         : 44 usecs
Worst two-way channel delay        : 45 usecs
Average two-way channel delay      : 45 usecs
Best round-trip-time               : 1752 usecs
Worst round-trip-time              : 3243 usecs
Average round-trip-time            : 2498 usecs
Average forward delay variation    : 1 usecs
Average reverse delay variation    : 1 usecs

DM queries sent                    : 2
DM responses received              : 2
DM queries timeout                 : 0
DM responses dropped due to errors : 0

```

monitor mpls lsp delay rsvp count detail

```

user@host> monitor mpls lsp delay rsvp LSP-A count 2 detail

(1)
Response code                      : Success
Querier transmit timestamp         : 1404129122 secs, 479955401 nsecs
Responder receive timestamp        : 1404129122 secs, 468519022 nsecs
Responder transmit timestamp       : 1404129122 secs, 470255123 nsecs
Querier receive timestamp          : 1404129122 secs, 481736403 nsecs
Current two-way channel delay      : 44 usecs
Current round-trip-time            : 1781 usecs
(2)
Response code                      : Success
Querier transmit timestamp         : 1404129123 secs, 480926210 nsecs
Responder receive timestamp        : 1404129123 secs, 469488696 nsecs
Responder transmit timestamp       : 1404129123 secs, 471130706 nsecs
Querier receive timestamp          : 1404129123 secs, 482613911 nsecs
Current two-way channel delay      : 45 usecs

```

```

Current round-trip-time           : 1687 usecs

Best two-way channel delay       : 44 usecs
Worst two-way channel delay      : 45 usecs
Average two-way channel delay    : 45 usecs
Best round-trip-time            : 1687 usecs
Worst round-trip-time           : 1781 usecs
Average round-trip-time         : 1734 usecs
Average forward delay variation  : 1 usecs
Average reverse delay variation  : 1 usecs

DM queries sent                  : 2
DM responses received            : 2
DM queries timedout              : 0
DM responses dropped due to errors : 0
regress@pro0-a> monitor mpls loss-delay-measurement lsp LSP1_A_to_B count 2
(1)
Current forward loss             : 0 packets
Current forward loss ratio       : 0.000000
Current forward throughput       : 0.957 kpps
Current reverse loss             : 0 packets
Current reverse loss ratio       : 0.000000
Current reverse throughput       : 0.962 kpps
Current two-way channel delay    : 48 usecs
Current round-trip-time         : 3476 usecs
(2)
Current forward loss             : 0 packets
Current forward loss ratio       : 0.000000
Current forward throughput       : 0.599 kpps
Current reverse loss             : 0 packets
Current reverse loss ratio       : 0.000000
Current reverse throughput       : 0.599 kpps
Current two-way channel delay    : 50 usecs
Current round-trip-time         : 1856 usecs

Cumulative forward transmit count : 1557
Cumulative forward loss           : 0 packets
Average forward loss ratio        : 0.000000
Average forward throughput        : 0.778 kpps
Cumulative reverse transmit count : 1562
Cumulative reverse loss           : 0 packets
Average reverse loss ratio        : 0.000000
Average reverse throughput        : 0.780 kpps

Best two-way channel delay       : 48 usecs
Worst two-way channel delay      : 50 usecs
Average two-way channel delay    : 49 usecs
Best round-trip-time            : 1856 usecs
Worst round-trip-time           : 3476 usecs
Average round-trip-time         : 2445 usecs
Average forward delay variation  : 1 usecs
Average reverse delay variation  : 1 usecs

LDM queries sent                 : 3
LDM responses received           : 3
LDM queries timedout             : 0
LDM responses dropped due to errors : 0

```

monitor mpls loss rsvp

Syntax `monitor mpls loss rsvp lsp-name`
 `<detail>`
 `<bytes>`
 `<count count>`
 `<interval seconds>`
 `<traffic-class traffic-class>`

Release Information Command introduced in Junos OS Release 14.2.

Description Perform an on-demand loss measurement and display the measured values for associated bidirectional MPLS ultimate hop popping (UHP) point-to-point label-switched paths (LSPs).

Options *lsp-name*—Name of the associated bidirectional MPLS UHP LSP for which the loss measurement is performed.

detail—(Optional) Display detailed output of the LSP loss measurement.

bytes—(Optional) Specify the measurement quantity for the LSP loss measurement as bytes. By default, LSP loss is measured in packets.



NOTE: The byte count of a packet sent or received over a channel counts only the payload, including the total length of that packet and excluding the headers, labels, and framing of the channel itself.

count count—(Optional) Specify the number of loss measurements to be carried out for the MPLS UHP LSP. For LSP loss measurement, the number of queries sent is the specified count number plus one additional query, because the LSP loss is measured using successive messages.

Default: 10

Range: 1 through 1000000

interval seconds—(Optional) Specify in seconds the interval between two successive query messages.

Range: 1 through 255 seconds

traffic-class traffic-class—(Optional) Specify the traffic class and enable traffic-class-statistics for the LSP loss measurement.

Range: 0 though 7

Required Privilege Level view

Related Documentation

- [monitor mpls delay rsvp on page 1150](#)
- [monitor mpls loss-delay rsvp on page 1159](#)

- [Example: Configuring On-Demand Loss and Delay Measurement on page 439](#)

List of Sample Output [monitor mpls lsp loss rsvp count on page 1156](#)

[monitor mpls lsp loss rsvp detail on page 1157](#)

Output Fields [Table 20 on page 1151](#) describes the output fields for the **monitor mpls loss rsvp** command. Output fields are listed in the approximate order in which they appear.

Table 21: monitor mpls loss rsvp Output Fields

Field Name	Field Description	Level of Output
Current forward loss	Difference between the current forward transmit count and the current forward receive count.	All Levels
Current forward loss ratio	Total packet loss (current forward loss divided by current forward transmit count).	All Levels
Current forward throughput	Current forward transmit count divided by 1000.	All Levels
Current reverse loss	Difference between the current reverse transmit count and the current reverse receive count.	All Levels
Current reverse loss ratio	Total packet loss (current reverse loss divided by current reverse transmit count).	All Levels
Current reverse throughput	Current reverse transmit count divided by 1000.	All Levels
Cumulative forward transmit count	Cumulative forward transmit counter value at the time the loss measurement message was originated.	All Levels
Cumulative forward loss	Cumulative forward loss counter value at the time the loss measurement message was originated.	All Levels
Average forward loss ratio	Average packet loss (current forward loss divided by current forward transmit count).	All Levels
Average forward throughput	Average forward transmit count divided by 1000.	All Levels
Cumulative reverse transmit count	Cumulative reverse transmit counter value at the time the loss measurement message was originated.	All Levels
Cumulative reverse loss	Difference between the cumulative reverse transmit count and the cumulative reverse receive count.	All Levels
Average reverse loss ratio	Average packet loss (average reverse loss divided by average reverse transmit count).	All Levels
Average reverse throughput	Average reverse transmit count divided by 1000.	All Levels
LM queries sent	Number of queries sent for loss measurement.	All Levels
LM responses received	Number of responses received for loss measurement queries.	All Levels

Table 21: monitor mpls loss rsvp Output Fields (*continued*)

Field Name	Field Description	Level of Output
LM queries timedout	Number of timed out queries sent for loss measurement.	All Levels
LM responses dropped due to errors	Number of loss measurement responses dropped due to errors.	All Levels
Response code	Status of the messages used for loss measurement. Response code can be one of the following: <ul style="list-style-type: none"> • Success—Successful response code. • Failed—Failed response code. 	detail
Origin timestamp	Time and date the loss measurement message is originated without any specific format (NTP and PTP).	detail
Forward transmit count	Forward transmit counter value at the time the loss measurement message was originated.	detail
Forward receive count	Forward receive counter value at the time the loss measurement message was originated.	detail
Reverse transmit count	Reverse transmit counter value at the time the loss measurement message was originated.	detail
Reverse receive count	Reverse receive counter value at the time the loss measurement message was originated.	detail
Current forward transmit count	Difference between the current forward transit count and the previous forward transit count.	detail
Current forward receive count	Difference between the current forward receive count and the previous forward receive count.	detail
Current reverse transmit count	Difference between the current reverse transit count and the previous reverse transit count.	detail
Current reverse receive count	Difference between the current reverse receive count and the previous reverse receive count.	detail

Sample Output

monitor mpls lsp loss rsvp count

```
user@host> monitor mpls lsp loss rsvp count 2
```

```
(1)
Current forward loss           : 0 packets
Current forward loss ratio     : 0.000000
Current forward throughput     : 1.006 kpps
Current reverse loss           : 0 packets
Current reverse loss ratio     : 0.000000
Current reverse throughput     : 1.007 kpps
(2)
Current forward loss           : 0 packets
```

```

Current forward loss ratio          : 0.000000
Current forward throughput          : 0.559 kpps
Current reverse loss                : 0 packets
Current reverse loss ratio          : 0.000000
Current reverse throughput          : 0.562 kpps

Cumulative forward transmit count   : 1559
Cumulative forward loss             : 0 packets
Average forward loss ratio          : 0.000000
Average forward throughput          : 0.782 kpps
Cumulative reverse transmit count   : 1563
Cumulative reverse loss             : 0 packets
Average reverse loss ratio          : 0.000000
Average reverse throughput          : 0.784 kpps

LM queries sent                     : 3
LM responses received                : 3
LM queries timedout                 : 0
LM responses dropped due to errors   : 0

```

monitor mpls lsp loss rsvp detail

```

user@host> monitor mpls lsp loss rsvp detail
(0)
Response code                       : Success
Origin timestamp                    : 1404129082 secs, 905571890 nsecs
Forward transmit count              : 83040
Forward receive count               : 83040
Reverse transmit count              : 83100
Reverse receive count               : 83100
(1)
Response code                       : Success
Origin timestamp                    : 1404129083 secs, 905048410 nsecs
Forward transmit count              : 83841
Forward receive count               : 83841
Reverse transmit count              : 83904
Reverse receive count               : 83904
Current forward transmit count      : 801
Current forward receive count       : 801
Current forward loss                : 0 packets
Current forward loss ratio          : 0.000000
Current forward throughput          : 0.801 kpps
Current reverse transmit count      : 804
Current reverse receive count       : 804
Current reverse loss                : 0 packets
Current reverse loss ratio          : 0.000000
Current reverse throughput          : 0.804 kpps
(2)
Response code                       : Success
Origin timestamp                    : 1404129084 secs, 904828715 nsecs
Forward transmit count              : 84423
Forward receive count               : 84423
Reverse transmit count              : 84487
Reverse receive count               : 84487
Current forward transmit count      : 582
Current forward receive count       : 582
Current forward loss                : 0 packets
Current forward loss ratio          : 0.000000
Current forward throughput          : 0.582 kpps
Current reverse transmit count      : 583
Current reverse receive count       : 583

```

Current reverse loss	: 0 packets
Current reverse loss ratio	: 0.000000
Current reverse throughput	: 0.583 kpps
Cumulative forward transmit count	: 1383
Cumulative forward loss	: 0 packets
Average forward loss ratio	: 0.000000
Average forward throughput	: 0.692 kpps
Cumulative reverse transmit count	: 1387
Cumulative reverse loss	: 0 packets
Average reverse loss ratio	: 0.000000
Average reverse throughput	: 0.694 kpps
LM queries sent	: 3
LM responses received	: 3
LM queries timeout	: 0
LM responses dropped due to errors	: 0

monitor mpls loss-delay rsvp

Syntax `monitor mpls loss-delay rsvp lsp-name`
 `<detail>`
 `<bytes>`
 `<count count>`
 `<interval seconds>`
 `<padding-size padding-size>`
 `<traffic-class traffic-class>`

Release Information Command introduced in Junos OS Release 14.2.

Description Perform a simultaneous on-demand loss and delay measurement using combined loss and delay messages, and display the measured values for associated bidirectional MPLS ultimate hop popping (UHP) point-to-point label-switched paths (LSPs).

Options *lsp-name*—Name of the associated bidirectional MPLS UHP LSP for which the delay measurement is performed.

detail—(Optional) Display detailed output of the LSP delay measurement.

bytes—(Optional) Specify the measurement quantity for the LSP loss measurement as bytes. By default, LSP loss is measured in packets.



NOTE: The byte count of a packet sent or received over a channel counts only the payload, including the total length of that packet and excluding the headers, labels, and framing of the channel itself.

count *count*—(Optional) Specify the number of delay measurements to be carried out for the MPLS UHP LSP. For LSP delay measurement, the number of queries sent is the specified count number plus one additional query, because the LSP delay is measured using successive messages.

Default: 10

Range: 1 through 1000000

interval *seconds*—(Optional) Specify in seconds the interval between two successive query messages.

Range: 1 through 255 seconds

padding-size *padding-size*—(Optional) Specify the length of padding TLV to be included in the query message.

Range: 0 through 1500

traffic-class *traffic-class*—(Optional) Specify the traffic class for the LSP delay measurement. When the traffic-class value is not specified, the default traffic-class code-point of 111 is used.

Range: 0 through 7

Required Privilege Level view

Related Documentation

- [monitor mpls loss rsvp on page 1154](#)
- [monitor mpls delay rsvp on page 1150](#)
- [Example: Configuring On-Demand Loss and Delay Measurement on page 439](#)

List of Sample Output [monitor mpls loss-delay rsvp count on page 1160](#)
[monitor mpls loss-delay rsvp count detail on page 1161](#)

Output Fields For output field descriptions, see the [monitor mpls loss rsvp](#) and [monitor mpls delay rsvp](#) commands.

Sample Output

monitor mpls loss-delay rsvp count

```
user@host> monitor mpls loss-delay rsvp LSP-A count 2

(1)
Current forward loss           : 0 packets
Current forward loss ratio     : 0.000000
Current forward throughput     : 0.957 kpps
Current reverse loss           : 0 packets
Current reverse loss ratio     : 0.000000
Current reverse throughput     : 0.962 kpps
Current two-way channel delay  : 48 usecs
Current round-trip-time       : 3476 usecs

(2)
Current forward loss           : 0 packets
Current forward loss ratio     : 0.000000
Current forward throughput     : 0.599 kpps
Current reverse loss           : 0 packets
Current reverse loss ratio     : 0.000000
Current reverse throughput     : 0.599 kpps
Current two-way channel delay  : 50 usecs
Current round-trip-time       : 1856 usecs

Cumulative forward transmit count : 1557
Cumulative forward loss           : 0 packets
Average forward loss ratio       : 0.000000
Average forward throughput       : 0.778 kpps
Cumulative reverse transmit count : 1562
Cumulative reverse loss           : 0 packets
Average reverse loss ratio       : 0.000000
Average reverse throughput       : 0.780 kpps

Best two-way channel delay       : 48 usecs
Worst two-way channel delay      : 50 usecs
Average two-way channel delay    : 49 usecs
Best round-trip-time            : 1856 usecs
Worst round-trip-time           : 3476 usecs
Average round-trip-time         : 2445 usecs
Average forward delay variation  : 1 usecs
Average reverse delay variation  : 1 usecs

LDM queries sent                : 3
```

```

LDM responses received          : 3
LDM queries timedout           : 0
LDM responses dropped due to errors : 0

```

monitor mpls loss-delay rsvp count detail

```
user@host> monitor mpls loss-delay rsvp LSP-A count 2 detail
```

```

(0)
Response code                  : Success
Forward transmit count         : 142049
Forward receive count          : 142049
Reverse transmit count         : 142167
Reverse receive count          : 142167
Querier transmit timestamp     : 1404129161 secs, 554422723 nsecs
Responder receive timestamp    : 1404129161 secs, 542877570 nsecs
Responder transmit timestamp   : 1404129161 secs, 546004545 nsecs
Querier receive timestamp      : 1404129161 secs, 557599327 nsecs

(1)
Response code                  : Success
Forward transmit count         : 143049
Forward receive count          : 143049
Reverse transmit count         : 143168
Reverse receive count          : 143168
Current forward transmit count : 1000
Current forward receive count  : 1000
Current forward loss           : 0 packets
Current forward loss ratio     : 0.000000
Current forward throughput     : 1.000 kpps
Current reverse transmit count : 1001
Current reverse receive count  : 1001
Current reverse loss           : 0 packets
Current reverse loss ratio     : 0.000000
Current reverse throughput     : 1.001 kpps
Querier transmit timestamp     : 1404129162 secs, 554465742 nsecs
Responder receive timestamp    : 1404129162 secs, 542919166 nsecs
Responder transmit timestamp   : 1404129162 secs, 545812736 nsecs
Querier receive timestamp      : 1404129162 secs, 557409175 nsecs
Current two-way channel delay  : 49 usecs
Current round-trip-time       : 2943 usecs

(2)
Response code                  : Success
Forward transmit count         : 143677
Forward receive count          : 143677
Reverse transmit count         : 143799
Reverse receive count          : 143799
Current forward transmit count : 628
Current forward receive count  : 628
Current forward loss           : 0 packets
Current forward loss ratio     : 0.000000
Current forward throughput     : 0.627 kpps
Current reverse transmit count : 631
Current reverse receive count  : 631
Current reverse loss           : 0 packets
Current reverse loss ratio     : 0.000000
Current reverse throughput     : 0.630 kpps
Querier transmit timestamp     : 1404129163 secs, 556698575 nsecs
Responder receive timestamp    : 1404129163 secs, 545150128 nsecs
Responder transmit timestamp   : 1404129163 secs, 546918408 nsecs
Querier receive timestamp      : 1404129163 secs, 558515047 nsecs
Current two-way channel delay  : 48 usecs

```

Current round-trip-time	: 1816 usecs
Cumulative forward transmit count	: 1628
Cumulative forward loss	: 0 packets
Average forward loss ratio	: 0.000000
Average forward throughput	: 0.813 kpps
Cumulative reverse transmit count	: 1632
Cumulative reverse loss	: 0 packets
Average reverse loss ratio	: 0.000000
Average reverse throughput	: 0.815 kpps
Best two-way channel delay	: 48 usecs
Worst two-way channel delay	: 49 usecs
Average two-way channel delay	: 49 usecs
Best round-trip-time	: 1816 usecs
Worst round-trip-time	: 3176 usecs
Average round-trip-time	: 2645 usecs
Average forward delay variation	: 1 usecs
Average reverse delay variation	: 0 usecs
LDM queries sent	: 3
LDM responses received	: 3
LDM queries timedout	: 0
LDM responses dropped due to errors	: 0

ping mpls bgp

Syntax ping mpls bgp *fec*
 <bottom-label-ttl>
 <count *count*>
 <destination *address*>
 <detail>
 <exp *forwarding-class*>
 <instance *routing-instance-name*>
 <logical-system (all | *logical-system-name*)>
 <size *bytes*>
 <source *source-address*>
 <sweep>

Release Information Command introduced in Junos OS Release 11.1.

Description Check the operability of MPLS BGP-signaled label-switched path (LSP) connections. Press Ctrl+c to interrupt a **ping mpls bgp** command.



NOTE: The **ping mpls bgp *fec*** command only supports single paths.

- Options**
- bottom-label-ttl**—(Optional) Time-to-live (TTL) value for the bottom label in the label stack. The range of values is 1 through 255. The default value is **255**.
 - count *count***—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is **5**.
 - destination *address***—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.
 - detail**—(Optional) Display detailed information about the echo requests sent and received.
 - exp *forwarding-class***—(Optional) Value of the forwarding class for the MPLS ping packets.
 - fec***—Ping a BGP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.
 - instance *routing-instance-name***—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.
 - logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on the specified logical system.
 - size *bytes***—(Optional) Size of the LSP ping request packet (88 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only BGP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level network

List of Sample Output [ping mpls bgp fec count on page 1164](#)

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately. To display the error codes, use the **detail** option (for example, **ping mpls bgp 10.255.245.222 detail**).

Sample Output

[ping mpls bgp fec count](#)

```
user@host> ping mpls bgp 10.255.245.222 count 10
!!!xxx...x--- 1sping statistics ---10 packets transmitted, 3 packets received,
70% packet loss 4 packets received with error status, not counted as received.
```

ping mpls lsp-end-point

Syntax	<pre>ping mpls lsp-end-point <i>prefix-name</i> <count <i>count</i>> <destination <i>address</i>> <detail> <exp <i>forwarding-class</i>> <instance <i>routing-instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <size <i>bytes</i>> <source <i>source-address</i>> <sweep></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>The size and sweep options were introduced in Junos OS Release 9.6.</p> <p>The instance option was introduced in Junos OS Release 10.0.</p>
Description	<p>Check the operability of MPLS label-switched path (LSP) endpoint connections. Type Ctrl+c to interrupt a ping mpls command.</p>
Options	<p>count <i>count</i>—(Optional) Number of ping requests to send. If count is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.</p> <p>destination <i>address</i>—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.</p> <p>detail—(Optional) Display detailed information about the echo requests sent and received.</p> <p>exp <i>forwarding-class</i>—(Optional) Value of the forwarding class for the MPLS ping packets.</p> <p>instance <i>routing-instance-name</i>—(Optional) Ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP connection.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on the specified logical system.</p> <p>prefix-name—LDP forwarding equivalence class (FEC) prefix or RSVP LSP endpoint address.</p> <p>size <i>bytes</i>—(Optional) Size of the LSP ping request packet. If the endpoint is LDP-based, the minimum size of the packet is 88 bytes. If the endpoint is RSVP-based, the minimum size of the packet is 100 bytes. The maximum size in either case is 65468 bytes.</p> <p>source <i>source-address</i>—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).</p> <p>sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).</p>

Additional Information If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level network

List of Sample Output [ping mpls lsp-end-point detail on page 1166](#)

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

[ping mpls lsp-end-point detail](#)

```
user@host> ping mpls lsp-end-point 10.255.245.119 detail
Route to end point address is via LDP FEC
Request for seq 1, to interface 67, label 100032
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, label 100032
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 67, label 100032
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 67, label 100032
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 67, label 100032
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

request mpls container-lsp

Syntax	request mpls container-lsp <logical-system (all <i>logical-system-name</i>)> <name <i>lsp-name</i> > <adjust-autobandwidth> <normalization>
Release Information	Command introduced in Junos OS Release 14.2.
Description	Manually trigger a bandwidth allocation adjustment for the container label-switched path (LSP).
Options	<p>none—Manually trigger a bandwidth allocation adjustment for all active member LSP paths.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>name <i>lsp-name</i>—(Optional) Manually trigger a bandwidth allocation adjustment on the specified member LSP only.</p> <p>adjust-autobandwidth—(Optional) Request LSP autobandwidth adjustment.</p> <p>normalization—(Optional) Request container LSP normalization.</p>
Required Privilege Level	clear, maintenance
Related Documentation	<ul style="list-style-type: none"> • show mpls container-lsp on page 1190 • clear mpls container-lsp on page 1147
List of Sample Output	request mpls container-lsp on page 1167 request mpls container-lsp on page 1167
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request mpls container-lsp

```
user@host> request mpls container-lsp lsp-name normalize
```

request mpls container-lsp

```
user@host> request mpls container-lsp normalize bandwidth bps
```

request mpls lsp adjust-autobandwidth

List of Syntax	Syntax on page 1168 Syntax (EX and QFX Series Switches) on page 1168
Syntax	<pre>request mpls lsp adjust-autobandwidth <logical-system (all <i>logical-system-name</i>)> <name <i>lsp-name</i>></pre>
Syntax (EX and QFX Series Switches)	<pre>request mpls lsp adjust-autobandwidth <name <i>lsp-name</i>></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	<p>Manually trigger a bandwidth allocation adjustment for active label-switched paths (LSPs).</p> <p>Without running this command, the bandwidth adjustment is recomputed at a configurable interval. The default interval is 5 minutes. If you do not want to wait for the periodic adjustment (for example, during a software demonstration), this command is useful.</p> <p>During bandwidth allocation adjustment, the LSP stays up to enable the bandwidth to be changed without dropping any traffic. This functionality is often referred to as <i>make-before-break</i>.</p>
Options	<p>none—Manually trigger a bandwidth allocation adjustment for all active LSP paths.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>name <i>lsp-name</i>—(Optional) Manually trigger a bandwidth allocation adjustment on the specified LSP only.</p>
Additional Information	<p>For this command to work properly, the following conditions must exist:</p> <ul style="list-style-type: none">• Automatic bandwidth allocation must be enabled on the LSP. The parameters for adjustment interval and maximum average bandwidth are not reset after you issue the request mpls lsp adjust-autobandwidth command.• The difference between the adjusted bandwidth and the current LSP path bandwidth must be greater than the threshold limit.
Required Privilege Level	clear, maintenance
Related Documentation	<ul style="list-style-type: none">• auto-bandwidth on page 838• Configuring Automatic Bandwidth Allocation for LSPs on page 257

List of Sample Output [request mpls lsp adjust-auto-bandwidth on page 1169](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[request mpls lsp adjust-auto-bandwidth](#)

```
user@host> request mpls lsp adjust-auto-bandwidth
```

show connections

List of Syntax [Syntax on page 1170](#)
[Syntax \(EX Series Switches\) on page 1170](#)

Syntax show connections
 <brief | extensive>
 <all | interface-switch | lsp-switch | p2mp-receive-switch | p2mp-transmit-switch |
 remote-interface-switch>
 <down | up | up-down>
 <history>
 <labels>
 <logical-system (all | *logical-system-name*)>
 <name>
 <status>

Syntax (EX Series Switches) show connections
 <brief | extensive>
 <all | interface-switch | lsp-switch | p2mp-receive-switch | p2mp-transmit-switch |
 remote-interface-switch>
 <down | up | up-down>
 <history>
 <labels>
 <name>
 <status>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.5 for EX Series switches.

Description Display information about the configured circuit cross-connect (CCC) connections.

Options **none**—Display the standard level of output for all configured CCC connections.

all—(Optional) Display all connections.

brief | extensive—(Optional) Display the specified level of output. Use history to display information about connection history. Use labels to display labels used for transmit and receive LSPs. Use status to display information about the connection and interface status.

interface-switch—(Optional) Display interface switch connections only.

lsp-switch—(Optional) Display LSP switch connections only.

p2mp-receive-switch—(Optional) Display point-to-multipoint LSP to local interfaces switch connections only.

p2mp-transmit-switch—(Optional) Display local interface to point-to-multipoint LSP switch connections only.

remote-interface-switch—(Optional) Display remote interface switch connections only.

down | up | up-down—(Optional) Display nonoperational, operational, or both kinds of connections.

history—(Optional) Display information about connection history.

labels—(Optional) Display labels used for transmit and receive.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display information about the specified connection only.

status—(Optional) Display information about the connection and interface status.

Required Privilege Level view

Output Fields [Table 22 on page 1171](#) describes the output fields for the **show connections** command. Output fields are listed in the approximate order in which they appear.

Table 22: show connections Output Fields

Field Name	Field Description
CCC and TCC connections [Link Monitoring On Off]	Whether link monitoring is enabled: On or Off .
Legend for Status (St)	Connection or circuit status. See the output's legend for an explanation of the status field values.
Legend for connection types	Type of connection: <ul style="list-style-type: none"> if-sw—Layer 2 switching cross-connect. rmt-if—Remote interface switch. While graceful restart is in progress, rmt-if will display a state (St) of Restart. lsp-sw—LSP stitching cross-connect. While graceful restart is in progress, lsp-sw will display a state (St) of Restart.
Legend for circuit types	Type of circuits: <ul style="list-style-type: none"> intf—Interface circuit. tlsp—Transmit LSP circuit. rlsp—Receive LSP circuit.
Connection/Circuit	Name of the configured CCC connection.
Type	Type of connection.
St	State of the connection.
Time last up	Time that the connection or circuit last transitioned to the Up (operational) state.

Table 22: show connections Output Fields (*continued*)

Field Name	Field Description
# Up trans	Number of times that the connection or circuit has transitioned to the Up (operational) state.

Sample Output

show connections

```

user@switch> show connections
CCC and TCC connections [Link Monitoring On]
  Legend for status (St)           Legend for connection types
  UN -- uninitialized             if-sw: interface switching
  NP -- not present               rmt-if: remote interface switching
  WE -- wrong encapsulation       lsp-sw: LSP switching
  DS -- disabled
  Dn -- down
  -> -- only outbound conn is up  Legend for circuit types
  <- -- only inbound conn is up   intf -- interface
  Up -- operational               tlsp -- transmit LSP
  RmtDn -- remote CCC down        rlsp -- receive LSP
  Restart -- restarting

CCC Graceful restart : Restarting

Connection/Circuit      Type   St      Time last up    # Up trans
IFSW-ed                 if-sw  Up       Aug  5 15:39:15      1
  so-1/0/2.0             intf   Up
  t1-0/1/2.0             intf   Up
SW-db                   rmt-if Restart      0
  so-1/0/3.0             intf   Up
  pro4-ca                 tlsp   Dn
  pro4-ac                 rlsp   NP

```

show link-management

Syntax	show link-management
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display Multiprotocol Label Switching (MPLS) peer and traffic engineering link information.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show link-management peer on page 1177 • show link-management routing on page 1179 • show link-management statistics on page 1182 • show link-management te-link on page 1184
List of Sample Output	show link-management on page 1176
Output Fields	Table 23 on page 1173 describes the output fields for the show link-management command. Output fields are listed in the approximate order in which they appear.

Table 23: show link-management Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is 0 through 64,000.
State	State of the peer: Up or Down .
Control address	Address to which a control channel is established.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296.
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296.
State	State of the control channel: Up or Down .
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is 1 through 4,294,967,295.
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is 0 through 4,294,967,295.

Table 23: show link-management Output Fields (*continued*)

Field Name	Field Description
Flags	Code that provides information about the control channel. Currently supports only code value R , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
TE links	Traffic-engineered links that are managed by their peer.
TE link name	Name of the traffic-engineered link.
State	State of the traffic-engineered link: Up , Down , or Init .
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET , Ethernet , Packet , and PDH .
Switching	Type of switching that can be performed on the traffic-engineered link. Supported values are PSC-1 and Packet .
Minimum bandwidth	Smallest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link (in bps).
Maximum bandwidth	Largest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
Total bandwidth	Sum of the bandwidth, in bits per second (bps) and megabits per second (Mbps), of all interfaces that are members of the link.
Available bandwidth	Sum of the bandwidths of all interfaces that are members of the link and that are not yet allocated (in bps).
Name	Name of the interface.
State	State of the interface: Up or Down .
Local ID	Identifier of the local side of the interface.
Remote ID	Identifier of the remote side of the interface.
Bandwidth	Bandwidth, in bps or Mbps, of the member interface.
Used	Whether the resource is allocated to an LSP: Yes or No .

Table 23: show link-management Output Fields *(continued)*

Field Name	Field Description
LSP-name	LSP name.

Sample Output

show link-management

```
user@host> show link-management
Peer name: PEER-A, System identifier: 11973
State: Up, Control address: 10.255.245.4
  CC local ID CC remote ID State      TxSeqNum  RcvSeqNum  Flags
    24547      24547 Up          1027      1026
TE links:
  pro4-ba

TE link name: pro4-ba, State: Init
Local identifier: 2662, Remote identifier: 0, Encoding: SDH/SONET, Switching:
PSC-1,
Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:
155.52Mbps,
Available bandwidth: 155.52Mbps
  Name      State Local ID Remote ID      Bandwidth Used  LSP-name
  so-1/0/2   Up      21271      0      155.52Mbps    No
```

show link-management peer

Syntax	show link-management peer <name <i>peer-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display Multiprotocol Label Switching (MPLS) peer link information.
Options	none —Display all peer link information. name <i>peer-name</i> —(Optional) Display information for the specified peer only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show link-management on page 1173 • show link-management routing on page 1179 • show link-management statistics on page 1182 • show link-management te-link on page 1184
List of Sample Output	show link-management peer on page 1178
Output Fields	Table 24 on page 1177 describes the output fields for the show link-management peer command. Output fields are listed in the approximate order in which they appear.

Table 24: show link-management peer Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is 0 through 64,000.
State	State of the peer: Up or Down .
Control address	Address to which a control channel is established.
Hello interval	How often the routing device sends Link Management Protocol (LMP) hello packets.
Hello dead interval	How long LMP waits before declaring the control channel to be dead. This is an interval during which the routing device receives no LMP hello packets from the neighbor on a control that is active or up.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296.
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296.

Table 24: show link-management peer Output Fields (*continued*)

Field Name	Field Description
State	State of the control channel: Up or Down .
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is 1 through 4,294,967,295 .
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is 0 through 4,294,967,295 .
Flags	Code that provides information about the control channel. Currently supports only code value R , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
TE links	Traffic-engineered links that are managed by their peer.

Sample Output

show link-management peer

```

user@host> show link-management peer
Peer name: sonet, System identifier: 41448
State: Up, Control address: 70.70.70.70
Hello interval: 10000, Hello dead interval: 30000
  CC local ID CC remote ID State      TxSeqNum  RcvSeqNum  Flags
    3265           0 ConfSnd         1          0 R
TE links:
to-sonet

```


show link-management routing

Syntax	show link-management routing <peer <name <i>name</i> > te-link <name <i>name</i> >> <resource <name <i>name</i> >>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display Multiprotocol Label Switching (MPLS) peer or traffic engineering link information from the routing process.
Options	<p>none—Display all peer and traffic-engineered link information.</p> <p>peer <name <i>name</i>>—(Optional) Display information for all peers or for the specified peer only.</p> <p>resource <name <i>name</i>>—(Optional) Display information for all resources or for the specified resource only.</p> <p>te-link <name <i>name</i>>—(Optional) Display information for all traffic-engineered forwarding paths or for the specified path only.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show link-management on page 1173 • show link-management peer on page 1177 • show link-management statistics on page 1182 • show link-management te-link on page 1184
List of Sample Output	show link-management routing on page 1181
Output Fields	Table 25 on page 1179 describes the output fields for the show link-management routing command. Output fields are listed in the approximate order in which they appear.

Table 25: show link-management routing Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is 0 through 64,000.
State	State of the peer: Up or Down.
Control address	Address to which a control channel is established.
Control channel	Interface over which control packets are sent.

Table 25: show link-management routing Output Fields (*continued*)

Field Name	Field Description
State	State of the control channel.
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: Up or Down .
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET , Ethernet , and Packet .
Minimum bandwidth	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
Maximum bandwidth	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
Total bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link.
Available bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
Resource	Forwarding adjacency LSP information.
Type	Type of resource. The type is always a forwarding adjacency LSP.
State	State of the LSP: Up or Down .
System Identifier	Internal identifier for the peer. The range of values is 0 through 64,000 .
Total bandwidth	Bandwidth resource, in bps or Mbps, on the TE-link learned from the routing process.
Traffic parameters	<ul style="list-style-type: none"> • Encoding—Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET, Ethernet, and Packet. • Switching—Type of switching that can be performed on the traffic-engineered link: PSC-1 and Packet. • Granularity—Layer 2 data for switching Layer 2 LSPs for this resource. Not supported. This value is always unknown.

Sample Output

show link-management routing

```

user@host> show link-management routing
Peer name: __rpd:fe-0/1/0.0, System identifier: 2147483649
State: Up, Control address: (null)
Control-channel          State
fe-0/1/0.0               Active

Peer name: __rpd:fe-0/1/2.0, System identifier: 2147483650
State: Up, Control address: (null)
Control-channel          State
fe-0/1/2.0               Active

Peer name: __rpd:so-0/2/0.0, System identifier: 2147483651
State: Down, Control address: (null)
Control-channel          State
so-0/2/0.0               State

Peer name: __rpd:so-0/2/1.0, System identifier: 2147483652
State: Down, Control address: (null)
Control-channel          State
so-0/2/1.0               State

...

TE link name: __rpd:fe-0/1/0.0, State: Up
Local identifier: 2147483649, Remote identifier: 0,
Local address: 192.168.37.66, Remote address: 192.168.37.66,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:fe-0/1/2.0, State: Up
Local identifier: 2147483650, Remote identifier: 0,
Local address: 192.168.37.73, Remote address: 192.168.37.73,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:so-0/2/0.0, State: Down
Local identifier: 2147483651, Remote identifier: 0,
Local address: 192.168.37.82, Remote address: 192.168.37.95,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 155.52Mbps,
Total bandwidth: 155.52Mbps, Available bandwidth: 155.52Mbps

...

Resource: falsp-bd, Type: LSP, State: Dn System identifier: 2147483652,
Total bandwidth: 0bps, Traffic parameters: Encoding: Packet, Switching: Packet,
Granularity: Unknown

Resource: falsp-be, Type: LSP, State: Up System identifier: 2147483654,
Total bandwidth: bw[1]=10Mbps, Traffic parameters: Encoding: Packet,
Switching: Packet, Granularity: Unknown

```

show link-management statistics

Syntax	<code>show link-management statistics</code> <code><peer <name <i>name</i>>></code>
Release Information	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display statistical information for Link Management Protocol (LMP) packets.
Options	none —Display information for all peers. peer <name <i>name</i>> —(Optional) Display information for all peers or for the specified peer only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show link-management on page 1173 • show link-management peer on page 1177 • show link-management routing on page 1179 • show link-management te-link on page 1184
List of Sample Output	show link-management statistics on page 1183
Output Fields	Table 26 on page 1182 describes the output fields for the show link-management statistics command. Output fields are listed in the approximate order in which they appear.

Table 26: show link-management statistics Output Fields

Field Name	Field Description
Received packets	Number of received packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Received bad packets	Number of received bad packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Small packets	Number of packets that are too small.
Wrong protocol version	Number of packets specifying the wrong LMP version.
Messages for unknown peer	Number of packets destined for an unknown peer.
Messages for bad state	Number of packets indicating a state that does not match the recipient.
Stale acknowledgments	Number of configAck and LinkSummaryAck packets received that have a stale message ID.

Table 26: show link-management statistics Output Fields (*continued*)

Field Name	Field Description
Stale negative acknowledgments	Number of configNack and LinkSummaryNack packets received that have a stale message ID.
Sent packets	Number of sent packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Retransmitted packets	Number of retransmitted packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Dropped packets	Number of packets sent, by message type, that have been dropped by the receiver after the LMP retransmission interval has been exceeded. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.

Sample Output

show link-management statistics

```

user@host> show link-management statistics peer pro4-a
Statistics for peer pro4-a
  Received packets
    Config: 1
    Hello: 2572
  Small packets: 0
  Wrong protocol version: 0
  Messages for unknown peer: 0
  Messages for bad state: 0
  Stale acknowledgments: 0
  Stale negative acknowledgments: 0
  Sent packets
    Config: 2
    ConfigAck: 1
    Hello: 2572
  Retransmitted packets
    Config: 1

```

show link-management te-link

Syntax	<code>show link-management te-link</code> <code><brief detail></code> <code><name <i>name</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.
Options	none —Display information for all traffic-engineered links. brief detail —(Optional) Display the specified level of output. name <i>name</i> —(Optional) Display information for the specified traffic-engineered link only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show link-management on page 1173 • show link-management peer on page 1177 • show link-management routing on page 1179 • show link-management statistics on page 1182
List of Sample Output	show link-management te-link on page 1185
Output Fields	Table 27 on page 1184 describes the output fields for the show link-management te-link command. Output fields are listed in the approximate order in which they appear.

Table 27: show link-management te-link Output Fields

Field Name	Field Description
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: Up or Down .
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include SDH/SONET , Ethernet , Packet , and PDH .

Table 27: show link-management te-link Output Fields (*continued*)

Field Name	Field Description
Switching	Type of switching that can be performed on the traffic-engineered link. Supported values are PSC-1 and Packet .
Minimum bandwidth	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
Maximum bandwidth	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link.
Total bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link (in bps).
Available Bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
Name	Name of the interface.
State	State of the interface: Up or Down .
Local ID	Identifier of the local side of the interface.
Remote ID	Identifier of the remote side of the interface.
Bandwidth	Bandwidth, in bps or Mbps, of the member interface.
Used	Whether the resource is allocated to an LSP: Yes or No .
LSP-name	LSP name.

Sample Output

show link-management te-link

```

user@host> show link-management te-link
TE link name: FA-bd, State: Up
  Local identifier: 4144, Remote identifier: 0, Local address: 2.2.2.1,
  Remote address: 2.2.2.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 0bps, Total bandwidth: 0bps,
  Available bandwidth: 0bps
    Name      State Local ID Remote ID      Bandwidth Used  LSP-name
    falsp-bd  Dn      43077      0             0bps No
TE link name: FA-be, State: Up
  Local identifier: 4145, Remote identifier: 0, Local address: 1.1.1.1,
  Remote address: 1.1.1.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 10Mbps, Total bandwidth: 10Mbps,
  Available bandwidth: 8Mbps
    Name      State Local ID Remote ID      Bandwidth Used  LSP-name
    falsp-be  Up      43076      0          10Mbps Yes  e2elasp-bf

```

show mpls admin-groups

List of Syntax	Syntax on page 1186 Syntax (EX Series Switches) on page 1186
Syntax	<pre>show mpls admin-groups <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	show mpls admin-groups
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>instance <i>instance-name</i> option added in Junos OS Release 15.1.</p>
Description	Display information about configured Multiprotocol Label Switching (MPLS) administrative groups.
Options	<p>none—Display information about the configured MPLS administrative groups.</p> <p>instance <i>instance-name</i>—(Optional) Display MPLS administrative group information for the specified instance. If <i>instance-name</i> is omitted, MPLS administrative group information for the master instance is displayed.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show mpls admin-groups on page 1186
Output Fields	<p>Table 28 on page 1186 describes the output fields for the show mpls admin-groups command. Output fields are listed in the approximate order in which they appear.</p>

Table 28: show mpls admin-groups Output Fields

Field Name	Field Description
Group	Name of the administrative group.
Bit index	Value assigned to the administrative group.

Sample Output

show mpls admin-groups

```
user@host> show mpls admin-groups
Group      Bit index
black      3
blue       2
```


gold	1
green	0

show mpls call-admission-control

List of Syntax	Syntax on page 1188 Syntax (EX Series Switches) on page 1188
Syntax	<pre>show mpls call-admission-control <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <lsp-name></pre>
Syntax (EX Series Switches)	<pre>show mpls call-admission-control <lsp-name></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>instance <i>instance-name</i> option added in Junos OS Release 15.1.</p>
Description	Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information.
Options	<p>none—Display CAC information for all LSPs.</p> <p>instance <i>instance-name</i>—(Optional) Display MPLS LSP CAC information for the specified instance. If instance-name is omitted, MPLS LSP CAC information for the master instance is displayed.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>lsp-name—(Optional) Display CAC information for the specified LSP only.</p>
Additional Information	The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by any Layer 2 connection at that class type.
Required Privilege Level	view
List of Sample Output	show mpls call-admission-control on page 1189
Output Fields	Table 29 on page 1188 describes the output fields for the show mpls call-admission-control command. Output fields are listed in the approximate order in which they appear.

Table 29: show mpls call-admission-control Output Fields

Field Name	Field Description
Available bandwidth	Current available bandwidth on each LSP path. Depending on whether the LSP is an E-LSP or a regular LSP, either per-class bandwidth or a single bandwidth value (corresponding to best-effort bandwidth at ct0) is displayed. The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by some Layer 2 connections at that class type.

Table 29: show mpls call-admission-control Output Fields (*continued*)

Field Name	Field Description
Layer2 connections	Different Layer 2 connections that had some bandwidth requirement and were admitted into an LSP path.
LSP name	LSP pathname.
Neighbor address	Neighbor address from which CAC and bandwidth booking are configured for Layer 2 circuits.
Circuit	Interface name and circuit information.
Primary	LSP's primary standby path.
Standby	LSP's secondary standby path.
VC bandwidth	Bandwidth constraints associated with a Layer 2 circuit route.

Sample Output

show mpls call-admission-control

```
user@host# show mpls call-admission-control
```

```

LSP name: pro1-be
*Primary
  Available bandwidth: 0bps

LSP name: pro1-be-1
*Primary
  Available bandwidth: 60kbps

LSP name: pro1-be-gold
*Primary
  Available bandwidth: <ct0 50kbps> <ct1 20kbps> <ct2 30kbps> <ct3 0bps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.0(vc 5)
    VC bandwidth: <ct0 50kbps> <ct1 40kbps> <ct2 40kbps>

LSP name: pro1-be-gold-2
*Primary
  Available bandwidth: <ct0 0bps> <ct1 40kbps> <ct2 40kbps> <ct3 0bps>

LSP name: pro1-be-silver
*Primary  prim1
  Available bandwidth: <ct0 10kbps> <ct1 20kbps> <ct2 0bps> <ct3 40kbps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>
  Standby  sec1
  Available bandwidth: <ct0 10kbps> <ct1 10kbps> <ct2 20kbps> <ct3 0bps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>

```

show mpls container-lsp

Syntax `show mpls container-lsp`
 `<brief | detail | extensive | terse>`
 `<count-active-routes>`
 `<defaults>`
 `<descriptions>`
 `<down | up>`
 `<egress>`
 `<ingress>`
 `<logical-system (all | logical-system-name)>`
 `<name name>`
 `<statistics>`
 `<transit>`
 `<unidirectional>`

Release Information Command introduced in Junos OS Release 14.2.

Description Display information about configured and active Multiprotocol Label Switching (MPLS) container label-switched paths (LSPs).

Options **none**—Display standard information about all configured and active member LSPs of the container LSP.

brief | detail | extensive | terse—(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

count-active-routes—(Optional) Show active routes for the container LSP.

defaults—(Optional) Display the default settings of the container LSP.

descriptions—(Optional) Display the container LSP descriptions. To view this information, you must configure the description statement at the **[edit protocol mpls container-lsp]** hierarchy level. Only the LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

egress | ingress—(Optional) Display the member LSPs ending at this routing device or originating from this routing device, respectively.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name *name*—(Optional) Display information about the specified LSP or group of LSPs.

statistics—(Optional) Display accounting information about LSPs. Statistics are not available for LSPs on the egress routing device, because the penultimate routing device in the LSP sets the label to 0. Also, as the packet arrives at the egress routing

device, the hardware removes its MPLS header and the packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.

transit—(Optional) Display LSPs transiting this routing device.

unidirectional—(Optional) Display unidirectional LSP information.

Required Privilege Level view

Related Documentation

- [request mpls container-lsp on page 1167](#)
- [clear mpls container-lsp on page 1147](#)

List of Sample Output [show mpls container-lsp on page 1195](#)
[show mpls container-lsp extensive on page 1195](#)

Output Fields [Table 30 on page 1191](#) describes the output fields for the **show mpls container-lsp** command. Output fields are listed in the approximate order in which they appear.

Table 30: show mpls container-lsp Output Fields

Field Name	Field Description	Level of Output
Ingress LSP	Information about the member LSPs on the ingress routing device. Each LSP has one line of output.	All levels
Container LSP name	Name of the container LSP.	All levels
Member LSP count	Number of member LSPs in the container LSP.	All levels
To	Destination (egress routing device) of the session.	brief
From	Source (ingress routing device) of the session.	brief detail
State	State of the LSP handled by this RSVP session: <ul style="list-style-type: none"> • Up • Dn (down) • Restart 	brief detail
Rt	Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	brief
P	Path. An asterisk (*) underneath this column indicates that the LSP is a primary path.	brief
ActivePath	(Ingress LSP) Name of the active path: Primary or Secondary.	detail extensive
LSPname	Name of the member LSP.	brief detail

Table 30: show mpls container-lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Egress LSP	Information about the LSPs on the egress routing device. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. Each session has one line of output.	All levels
Transit LSP	Number of LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
Min LSPs	Minimum number of member LSPs. Default: 1	extensive
Max LSPs	Number of member LSPs that the container LSP can have at maximum. Default: 64 (due to ECMP limit)	extensive
Aggregate bandwidth	Sum of the bandwidths of all member LSPs.	extensive
NormalizeTimer	Duration between two normalization events. When not configured, 21600 seconds (6 hours) is set as the default value.	extensive
NormalizeThreshold	Change in aggregate LSP utilization to trigger splitting or merging expressed in percentage.	extensive
Max Signaling BW	Maximum bandwidth used to signal LSPs after a normalization event. Default value is 0 bps. When not configured, the value is inherited from the splitting bandwidth configuration. NOTE: Between two normalization events, when auto-bandwidth adjustment happens, the per-LSP auto-bandwidth configuration and thresholds are used, instead of the maximum signaling bandwidth threshold.	extensive
Min Signaling BW	Minimum bandwidth used to signal LSPs after a normalization event. Default value is 0 bps. When not configured, the value is inherited from the merging bandwidth configuration. NOTE: Between two normalization events, when auto-bandwidth adjustment happens, the per-LSP auto-bandwidth configuration and thresholds are used, instead of the minimum signaling bandwidth threshold.	extensive
Splitting BW	Bandwidth used for LSP splitting and merging. Default value is 0 bps. When not configured, the value is inherited from the auto-bandwidth maximum bandwidth configuration.	extensive
Merging BW	Bandwidth used for LSP splitting and merging. Default value is 0 bps. When not configured, the value is inherited from the auto-bandwidth minimum bandwidth configuration.	extensive

Table 30: show mpls container-lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
LSPtype		extensive
LoadBalance		extensive
MinBW	Minimum LSP bandwidth in bps related to auto-bandwidth.	extensive
AdjustTimer	Total amount of time in seconds allowed before LSP bandwidth adjustment take place. Range: 300 through 315360000 seconds	extensive
Max AvgBW util	Current value of the actual maximum average bandwidth utilization in bps.	extensive
Overflow limit	Threshold overflow limit.	extensive
Underflow limit	Threshold underflow limit.	extensive
Encoding type		extensive
Switching type		extensive
GPID		extensive
Priorities	Setup priority and hold priority values. For setup priority, 0 and 7 is the highest and lowest priority, respectively. When not explicitly configured, 7 and 0 are set as the default values for the setup priority and hold priority, respectively.	extensive
Bandwidth		extensive
SmartOptimizeTimer	Time in seconds allowed before path reoptimization.	extensive
Computed ERO	Computed explicit route. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (S) or loose (L).	extensive

Table 30: show mpls container-lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Received RRO	<p>Received record route.</p> <p>RRO is a series of hops, each with an address followed by a flag. In most cases, the received RRO is the same as the computed ERO. If the received RRO is different from the computed ERO, there is a topology change in the network, and the route is taking a detour.</p> <p>The following flags identify the protection capability and status of the downstream node:</p> <ul style="list-style-type: none"> • 0x01—Local protection available. The link downstream from this node is protected by a local repair mechanism. This flag can be set only if the local protection flag was set in the SESSION_ATTRIBUTE object of the corresponding path message. • 0x02—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously). • 0x03—Combination of 0x01 and 0x02. • 0x04—Bandwidth protection. The downstream routing device has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section. • 0x08—Node protection. The downstream routing device has a backup path providing protection against link and node failure on the corresponding path section. If the downstream routing device can set up only a link-protection backup path, the local protection available bit is set but the node protection bit is cleared. • 0x09—Detour is established. Combination of 0x01 and 0x08. • 0x10—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engine LSP. This flag indicates to the ingress legacy edge router (LER) of this LSP that it should be rerouted. • 0x20—Node ID. Indicates that the address specified in the RRO's IPv4 or IPv6 sub-object is a node ID address, which refers to the router address or router ID. Nodes must use the same address consistently. • 0xb—Detour is in use. Combination of 0x01, 0x02, and 0x08. 	extensive
Make-before-break		extensive
Record Route		extensive
Automatic Autobw adjustment succeeded		extensive
CSPF		extensive
Created	Date and time the LSP was created.	extensive

Sample Output

show mpls container-lsp

```

user@host> show mpls container-lsp
Ingress LSP: 1 sessions
Container LSP name
test
To          From          State Rt P    Member LSP count
                ActivePath                2
10.255.107.76 10.255.107.78 Up    0 *    test-1
10.255.107.76 10.255.107.78 Up    0 *    test-2
Total 2 displayed, Up 2, Down 0

naling Egress LSP: 1 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show mpls container-lsp extensive

```

user@host> show mpls container-lsp extensive
Ingress LSP: 1 sessions
Container LSP name: test, Member count: 2
Normalization
Min LSPs: 2, Max LSPs: 64, Aggregate bandwidth: 0bps
NormalizeTimer: 1800 secs, NormalizeThreshold: 0%
Max Signaling BW: 2kbps, Min Signaling BW: 2kbps, Splitting BW: 5Mbps, Merging
BW: 2kbps
Normalization in 989 second(s)
10.255.107.76
From: 10.255.107.78, State: Up, ActiveRoute: 0, LSPname: test-1
ActivePath: (primary)
LSPtype: Dynamic Configured, Penultimate hop popping
LoadBalance: Random
Autobandwidth
MinBW: 1000bps
AdjustTimer: 300 secs
Max AvgBW util: 0bps, Bandwidth Adjustment in 89 second(s).
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up, No-decrement-ttl
Priorities: 7 0
Bandwidth: 1000bps
SmartOptimizeTimer: 180
Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
1.3.0.2 S 1.7.0.1 S
Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
1.3.0.2 1.7.0.1
11 Jul 13 20:08:26.613 Make-before-break: Switched to new instance
10 Jul 13 20:08:04.360 Record Route: 1.3.0.2 1.7.0.1
9 Jul 13 20:08:04.360 Up
8 Jul 13 20:08:04.360 Automatic Autobw adjustment succeeded: BW changes from
2000 bps to 1000 bps
7 Jul 13 20:08:04.314 Originate make-before-break call
6 Jul 13 20:08:04.314 CSPF: computation result accepted 1.3.0.2 1.7.0.1
5 Jul 13 20:05:02.423 Selected as active path
4 Jul 13 20:05:02.422 Record Route: 1.3.0.2 1.7.0.1
3 Jul 13 20:05:02.421 Up

```

```

    2 Jul 13 20:05:02.376 Originate Call
    1 Jul 13 20:05:02.376 CSPF: computation result accepted 1.3.0.2 1.7.0.1
Created: Sat Jul 13 20:03:03 2013
10.255.107.76
From: 10.255.107.78, State: Up, ActiveRoute: 0, LSPName: test-2
ActivePath: (primary)
LSPTYPE: Dynamic Configured, Penultimate hop popping
LoadBalance: Random
Autobandwidth
MinBW: 1000bps
AdjustTimer: 300 secs
Max AvgBW util: 0bps, Bandwidth Adjustment in 89 second(s).
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 0, Underflow Max AvgBW: 0bps
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary State: Up, No-decrement-ttl
  Priorities: 7 0
  Bandwidth: 1000bps
  SmartOptimizeTimer: 180
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
1.2.0.2 S 1.4.0.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    1.2.0.2 1.4.0.2
    11 Jul 13 20:08:05.363 Make-before-break: Switched to new instance
    10 Jul 13 20:08:04.450 Record Route: 1.2.0.2 1.4.0.2
    9 Jul 13 20:08:04.449 Up
    8 Jul 13 20:08:04.449 Automatic Autobw adjustment succeeded: BW changes from
2000 bps to 1000 bps
    7 Jul 13 20:08:04.327 Originate make-before-break call
    6 Jul 13 20:08:04.327 CSPF: computation result accepted 1.2.0.2 1.4.0.2
    5 Jul 13 20:05:00.849 Selected as active path
    4 Jul 13 20:05:00.841 Record Route: 1.3.0.2 1.7.0.1
    3 Jul 13 20:05:00.831 Up
    2 Jul 13 20:05:00.513 Originate Call
    1 Jul 13 20:05:00.502 CSPF: computation result accepted 1.3.0.2 1.7.0.1
Created: Sat Jul 13 20:03:03 2013
Total 2 displayed, Up 2, Down 0

Egress LSP: 1 sessions
Total 0 displayed, Up 0, Down 0
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show mpls context-identifier

Syntax	show mpls context-identifier <brief detail> <logical-system (all <i>logical-system-name</i>)> <primary>; <protector>;
Release Information	Command introduced in Junos OS Release 11.4R3.
Description	Display information about configured egress protection context identifiers.
Options	<p>none—Display standard information about egress protection.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>primary—(Optional) Perform this operation on the primary node.</p> <p>protector—(Optional) Perform this operation on the protector node.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP</i> • <i>Example: Configuring MPLS Egress Protection for Layer 3 VPN Services</i>
List of Sample Output	show mpls context-identifier detail (Protector) on page 1198 show mpls context-identifier detail (Primary) on page 1198
Output Fields	Table 31 on page 1197 describes the output fields for the show mpls egress-protection detail command. Output fields are listed in the approximate order in which they appear.

Table 31: show mpls lsp Output Fields

Field Name	Field Description	Level of Output
ID	Context identifier.	All levels
Type	Indicates node type: protector or primary	All levels
Metric	MPLS cost value of the context identifier route. This route appears in inet.0 on the protector and primary nodes. On the protector node, the metric is a larger number.	All levels
Mode	Indicates advertise-mode : proxy or alias	detail
Context table	Name of the MPLS routing table created for egress protection.	All levels

Table 31: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Context LSPs	Names of the LSPs that have egress protection configured. Loopback interface addresses of the devices from which the LSPs are originated.	detail
Total	Total number of primary and protector nodes.	All levels
Primary	Number of primary nodes.	All levels
Protector	Number of protector nodes.	All levels

Sample Output

show mpls context-identifier detail (Protector)

```

user@host> show mpls context-identifier detail
ID: 166.1.3.1
Type: protector, Metric: 16777215, Mode: alias
Context table: __166.1.3.1__.mpls.0, Label out: 299968

```

Sample Output

show mpls context-identifier detail (Primary)

```

user@host> show mpls context-identifier detail

ID: 166.1.3.1
Type: primary, Metric: 1, Mode: alias

Total 1, Primary 1, Protector 0

```

show mpls cspf

List of Syntax	Syntax on page 1199 Syntax (EX Series Switches) on page 1199
Syntax	<pre>show mpls cspf <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	show mpls cspf
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>instance <i>instance-name</i> option added in Junos OS Release 15.1.</p>
Description	Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics.
Options	<p>none—Display MPLS CSFP statistics.</p> <p>instance <i>instance-name</i>—(Optional) Display MPLS CSPF information for the specified instance. If <i>instance-name</i> is omitted, MPLS CSPF information for the master instance is displayed.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show mpls cspf on page 1200
Output Fields	Table 32 on page 1199 describes the output fields for the show mpls cspf command. Output fields are listed in the approximate order in which they appear.

Table 32: show mpls cspf Output Fields

Field Name	Field Description
Queue length	Number of LSPs queued for automatic path computation.
current	Current queue length.
maximum	Maximum queue length (high-water mark).
dequeued	Number of aborted computation attempts.
Paths	Counters for label-switched path computations.
total	Sum of the next four fields.

Table 32: show mpls cspf Output Fields (*continued*)

Field Name	Field Description
successful	Number of path computations that were successfully completed.
no route	Number of path computations that failed because the destination is unreachable.
Sys Error	Number of path computations that failed because of lack of memory.
CSPFs	Total number of CSPF computations. A single path might require multiple CSPF computations.
Time	Time, in seconds, required to perform the label-switched path computation.
Total	Total amount of time consumed by the CSPF path computation algorithm.
CSPFs	Total number of CSPF computations.
Avg per CSPF	Average amount of time required for each CSPF computation.
% of rpd	Percentage of routing process CPU used in the CSPF computation.

Sample Output

show mpls cspf

```

user@host> show mpls cspf
CSPF statistics
Queue length  current      maximum      dequeued
              0           0           0
Paths          total      successful      no route      sys error      CSPFs
              0           0           0           0           0
Time (secs)    total      CSPFs      avg per CSPF      % of rpd
              0.000000    0.000000    0.000000    0.0000

```

show mpls diffserv-te

List of Syntax	Syntax on page 1201 Syntax (EX Series Switches) on page 1201
Syntax	<pre>show mpls diffserve-te <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	show mpls diffserve-te
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>instance <i>instance-name</i> option added in Junos OS Release 15.1.</p>
Description	Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) Differentiated Services (DiffServ) class and preemption priority information.
Options	<p>none—Display DiffServ classes and priorities used by MPLS LSPs.</p> <p>instance <i>instance-name</i>—(Optional) Display DiffServ classes and priorities used by MPLS LSPs for the specified instance. If <i>instance-name</i> is omitted, DiffServ information for the master instance is displayed.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show mpls diffserv-te on page 1202
Output Fields	<p>Table 33 on page 1201 describes the output fields for the show mpls diffserv-te command. Output fields are listed in the approximate order in which they appear.</p>

Table 33: show mpls diffserv-te Output Fields

Field Name	Field Description
Bandwidth model	Bandwidth constraint model supported. The maximum allocation model (MAM) for EXP-inferred LSPs (E-LSPs) is currently supported.
TE class	DiffServ traffic engineering class.
Traffic class	<p>MPLS class type that corresponds to the DiffServ traffic engineering class:</p> <ul style="list-style-type: none"> • ct0—Best effort • ct1—Assured forwarding • ct2—Expedited forwarding • ct3—Network control

Table 33: show mpls diffserv-te Output Fields (*continued*)

Field Name	Field Description
Priority	MPLS preemption priority for this class type, a value from 0 through 7. Interior gateway protocols (IGPs) distribute information about the available bandwidth for each traffic engineering class.

Sample Output

show mpls diffserv-te

```
user@host> show mpls diffserv-te
Bandwidth model: Maximum Allocation Model with support for E-LSPs.
TE class      Traffic class  Priority
te0           ct0            3
te1           ct1            2
```


show mpls egress-protection


Syntax	show mpls egress-protection <brief detail> <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced in Junos OS Release 11.4R3.
Description	Display information about egress protection.
<div>  <p>NOTE: Use this command on the device configured as the protector PE router to display information about egress protection. If you use this command on the device configured as the primary PE router, no output is displayed.</p> </div>	
Options	<p>none—Display standard information about egress protection.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring MPLS Egress Protection for Layer 3 VPN Services</i> • <i>Example: Configuring Layer 3 VPN Egress Protection with RSVP and LDP</i>
List of Sample Output	show mpls egress-protection detail (Centralized Protector) on page 1204 show mpls egress-protection detail (Collocated Protector) on page 1204
Output Fields	Table 31 on page 1197 describes the output fields for the show mpls egress-protection detail command. Output fields are listed in the approximate order in which they appear.

Table 34: show mpls lsp Output Fields

Field Name	Field Description
Instance	Indicates egress instance name
Type	Indicates type of the VRF. It can be either local-vrf or remote-vrf
RIB	Indicates the edge-protection created routing table
Context-Id	Indicates the context-ID associated with the RIB.

Table 34: show mpls lsp Output Fields (*continued*)

Field Name	Field Description
Interface / EnhancedLookup	Show VT interfaces associated with the backup RIB. Shows Enhanced-lookup for MX Series 3D Universal Edge Routers with the Enhanced IP Network Services mode configured using the network-services enhanced-ip statement at the [edit chassis] hierarchy level.

Sample Output

show mpls egress-protection detail (Centralized Protector)

```

user@host> show mpls egress-protection detail

Instance           Type           Protection-Type
rsite1              remote-vrf     Protector
  RIB __99.99.1.4-rsite1__.inet.0, Context-Id 99.99.1.4, Enhanced-lookup
  Route Target 1:1
rsite24             remote-vrf     Protector
  RIB __99.99.1.4-rsite24__.inet.0, Context-Id 99.99.1.4, Enhanced-lookup
  Route Target 100:1023

```

Sample Output

show mpls egress-protection detail (Collocated Protector)

```

user@host> show mpls egress-protection detail

Instance           Type           Protection-Type
site2              local-vrf      Protector
  RIB __66.6.6.6-site2__.inet.0, Context-Id 66.6.6.6, Interface vt-1/3/0.87031809

  Route Target 100:251
site12             local-vrf      Protector
  RIB __66.6.6.6-site12__.inet.0, Context-Id 66.6.6.6, Interface vt-1/3/0.87031808

  Route Target 100:250
  Route Target 100:251

site2              local-vrf      Protector
  RIB __66.6.6.6-site2__.inet.0, Context-Id 66.6.6.6, Interface vt-1/3/0.87031809

  Route Target 100:251

```

show mpls interface

List of Syntax	Syntax on page 1205 Syntax (EX Series Switches) on page 1205
Syntax	show mpls interface <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show mpls interface
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. instance <i>instance-name</i> option added in Junos OS Release 15.1.
Description	Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces.
Options	none —Display information about MPLS-enabled interfaces. instance <i>instance-name</i> —(Optional) Display information about MPLS-enabled interfaces for the specified routing instance. If <i>instance-name</i> is omitted, information about MPLS-enabled interfaces is displayed for the master instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Additional Information	MPLS is enabled on an interface when the interface is configured with both the set protocol mpls interface <i>interface-name</i> and set interface <i>interface-name</i> unit 0 family mpls statements.
Required Privilege Level	view
List of Sample Output	show mpls interface on page 1206
Output Fields	Table 35 on page 1205 describes the output fields for the show mpls interface command. Output fields are listed in the approximate order in which they appear.

Table 35: show mpls interface Output Fields

Field Name	Field Description
Interface	Name of the interface.
State	State of the interface: Up or Dn (down).
Administrative groups	Administratively assigned colors of the link.

Table 35: show mpls interface Output Fields (*continued*)

Field Name	Field Description
Maximum labels	Maximum number of MPLS labels upon which MPLS can operate on a logical interface. This is configured using the maximum-labels statement at the [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] or the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] hierarchy levels.
Static protection revert time	Time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path. This is configured using the protection-revert-time statement at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels.
Always mark connection protection tlv	Enabled or Disabled: Enabled indicates that the always-mark-connection-protection-tlv statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. When this statement is configured, it marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, the switch-away-lsps statement must be configured.
Switch away lsps	Enabled or Disabled: Enabled indicates that the switch-away-lsps statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. This enables you to switch an LSP away from a network node using a bypass LSP. This feature can be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.

Sample Output

show mpls interface

```
user@host> show mpls interface
```

```
Interface: ge-0/2/1.57
State: Up
Administrative group: <none>
Maximum labels: 5
Static protection revert time: 5 seconds
Always mark connection protection tlv: Disabled
Switch away lsps : Disabled
```

show mpls label usage

Syntax `show mpls label usage`
`<logical-system (all | logical-system-name)>`

Release Information Command introduced in Junos OS Release 15.1.

Description Show the available label space resource in RPD and also the applications that use the label space in RPD. There are four different label spaces currently used in MPLS—namely LSI, Dynamic, Block and Static. Each label space has a fixed number and cannot grow beyond the fixed value. Using this command, the administrator can monitor the available labels in each label space and the applications that are using the labels. Based on the availability of labels, the administrator can decide to stop any service and free some labels or use other service where the labels are available.

Options **none**— Display the available labels in each label space and the applications using the labels.

logical-system (all | logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information Once the label space crosses the threshold, a new syslog message is added.

“<label-space-name> label space usage crossed threshold limit of 90%”.

For instance, LSI label space usage crossed threshold limit of 90%

Required Privilege Level view

List of Sample Output [show mpls label usage on page 1207](#)

Output Fields [Table 36 on page 1207](#) describes the output fields for the **show mpls label usage** command. Output fields are listed in the order in which they appear.

Table 36: show mpls label usage Fields

Field Name	Field Description
Label Space	The label spaces currently used in MPLS.
Available	Indicates the number of freely available labels and also the percentage of the label space available.
Applications	The applications that use the MPLS label spaces.

Sample Output

show mpls label usage

```
user@host> show mpls label usage
```

Label space	Available	Applications
LSI vrf-table-label	155331(75%)	BGP/LDP VPLS with no-tunnel-services, BGP L3VPN with
Dynamic	265645(50%)	RSVP, LDP, PW, L3VPN
Block	46189(30%)	BGP/LDP VPLS with tunnel-services, BGP L2VPN
Static	38575(67.12%)	Static LSP, Static PW

show mpls lsp

List of Syntax [Syntax on page 1209](#)
 [Syntax \(EX Series Switches\) on page 1209](#)

Syntax show mpls lsp
 <brief | detail | extensive | terse>
 <autobandwidth>
 <bidirectional | unidirectional>
 <bypass>
 <count-active-routes>
 <defaults>
 <descriptions>
 <down | up>
 <externally-controlled>
 <externally-provisioned>
 <logical-system (all | *logical-system-name*)>
 <lsp-type>
 <name *name*>
 <p2mp>
 <statistics>
 <transit>

Syntax (EX Series Switches) show mpls lsp
 <brief | detail | extensive | terse>
 <bidirectional | unidirectional>
 <bypass>
 <descriptions>
 <down | up>
 <externally-controlled>
 <externally-provisioned>
 <lsp-type>
 <name *name*>
 <p2mp>
 <statistics>
 <transit>

Release Information Command introduced before Junos OS Release 7.4.
 defaults option added in Junos OS Release 8.5.
 Command introduced in Junos OS Release 9.5 for EX Series switches.
 autobandwidth option added in Junos OS Release 11.4.
 externally-controlled option added in Junos OS Release 12.3.
 externally-provisioned option added in Junos OS Release 13.3.
 Command introduced in Junos OS Release 13.2X51-D15 for QFX Series.
 instance *instance-name* option added in Junos OS Release 15.1.

Description Display information about configured and active dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options **none**—Display standard information about all configured and active dynamic MPLS LSPs.

brief | detail | extensive | terse—(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

autobandwidth—(Optional) Display automatic bandwidth information. This option is explained separately (see [show mpls lsp autobandwidth](#)).

bidirectional | unidirectional—(Optional) Display bidirectional or unidirectional LSP information, respectively.

bypass—(Optional) Display LSPs used for protecting other LSPs.

count-active-routes—(Optional) Display active routes for LSPs.

defaults—(Optional) Display the MPLS LSP default settings.

descriptions—(Optional) Display the MPLS label-switched path (LSP) descriptions. To view this information, you must configure the description statement at the **[edit protocol mpls lsp]** hierarchy level. Only LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

externally-controlled—(Optional) Display the LSPs that are under the control of an external Path Computation Element (PCE).

externally-provisioned—(Optional) Display the LSPs that are generated dynamically and provisioned by an external Path Computation Element (PCE).

instance *instance-name*—(Optional) Display MPLS LSP information for the specified instance. If *instance-name* is omitted, MPLS LSP information is displayed for the master instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-type—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

name *name*—(Optional) Display information about the specified LSP or group of LSPs.

p2mp—(Optional) Display information about point-to-multipoint LSPs.

statistics—(Optional) (Ingress and transit routers only) Display accounting information about LSPs. Statistics are not available for LSPs on the egress routing device, because the penultimate routing device in the LSP sets the label to 0. Also, as the packet arrives at the egress routing device, the hardware removes its MPLS header and the

packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.



NOTE: If a bypass LSP is configured for the primary static LSP, display cumulative statistics of packets traversing through the protected LSP and bypass LSP when traffic is re-optimized when the protected LSP link is restored. (Bypass LSPs are not supported on QFX Series switches.)

When used with the **bypass** option (**show mpls lsp bypass statistics**), display statistics for the traffic that flows only through the bypass LSP.

transit—(Optional) Display LSPs transiting this routing device.

Required Privilege Level

view

Related Documentation

- [clear mpls lsp on page 1145](#)
- [show mpls lsp autobandwidth on page 1227](#)

List of Sample Output

[show mpls lsp defaults on page 1218](#)
[show mpls lsp descriptions on page 1218](#)
[show mpls lsp detail on page 1218](#)
[show mpls lsp detail \(When Egress Protection Is in Standby Mode\) on page 1219](#)
[show mpls lsp detail \(When Egress Protection Is in Effect During a Local Repair\) on page 1219](#)
[show mpls lsp extensive on page 1220](#)
[show mpls lsp ingress extensive on page 1222](#)
[show mpls lsp extensive \(automatic bandwidth adjustment enabled\) on page 1223](#)
[show mpls lsp bypass extensive on page 1224](#)
[show mpls lsp p2mp on page 1224](#)
[show mpls lsp p2mp detail on page 1224](#)
[show mpls lsp detail count-active-routes on page 1225](#)
[show mpls lsp statistics extensive on page 1226](#)

Output Fields

Table 37 on page 1211 describes the output fields for the **show mpls lsp** command. Output fields are listed in the approximate order in which they appear.

Table 37: show mpls lsp Output Fields

Field Name	Field Description	Level of Output
Ingress LSP	Information about LSPs on the ingress routing device. Each session has one line of output.	All levels
Egress LSP	Information about the LSPs on the egress routing device. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. Each session has one line of output.	All levels

Table 37: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transit LSP	Number of LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
P2MP name	Name of the point-to-multipoint LSP. Dynamically generated P2MP LSPs used for VPLS flooding use dynamically generated P2MP LSP names. The name uses the format <i>identifier:vpls:router-id:routing-instance-name</i> . The <i>identifier</i> is automatically generated by Junos OS.	All levels
P2MP branch count	Number of destination LSPs the point-to-multipoint LSP is transmitting to.	All levels
P	An asterisk (*) under this heading indicates that the LSP is a primary path.	All levels
address	(detail and extensive) Destination (egress routing device) of the LSP.	detail extensive
To	Destination (egress routing device) of the session.	brief
From	Source (ingress routing device) of the session.	brief detail
State	State of the LSP handled by this RSVP session: Up , Dn (down), or Restart .	brief detail
Active Route	Number of active routes (prefixes) installed in the forwarding table. For ingress LSPs, the forwarding table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (mpls.0).	detail extensive
Rt	Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	brief
P	Path. An asterisk (*) underneath this column indicates that the LSP is a primary path.	brief
ActivePath	(Ingress LSP) Name of the active path: Primary or Secondary .	detail extensive
LSPname	Name of the LSP.	brief detail
Statistics	Displays the number of packets and the number of bytes transmitted over the LSP. These counters are reset to zero whenever the LSP path is optimized (for example, during an automatic bandwidth allocation).	extensive
Aggregate statistics	Displays the number of packets and the number of bytes transmitted over the LSP. These counters continue to iterate even if the LSP path is optimized. You can reset these counters to zero using the clear mpls lsp statistics command.	extensive
Packets	Displays the number of packets transmitted over the LSP.	brief extensive
Bytes	Displays the number of bytes transmitted over the LSP.	brief extensive

Table 37: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
DiffServInfo	Type of LSP: multiclass LSP (multiclass diffServ-TE LSP) or Differentiated-Services-aware traffic engineering LSP (diffServ-TE LSP).	detail
LSPtype	Type of LSP: static Static configured or dynamic Dynamic configured . Also indicates if the LSP is a Penultimate hop popping LSP or an Ultimate hop popping LSP.	detail extensive
Bypass	(Bypass LSP) Destination address (egress routing device) for the bypass LSP.	All levels
LSPpath	Indicates whether the RSVP session is for the primary or secondary LSP path. LSPpath can be either primary or secondary and can be displayed on the ingress, egress, and transit routing devices.	detail
Bidir	(GMPLS) The LSP allows data to travel in both directions between GMPLS devices.	All levels
Bidirectional	(GMPLS) The LSP allows data to travel both ways between GMPLS devices.	All levels
FastReroute desired	Fast reroute has been requested by the ingress routing device.	detail
Link protection desired	detail	
Node/Link protection desired	Link protection has been requested by the ingress routing device.	detail extensive
LoadBalance	(Ingress LSP) CSPF load-balancing rule that was configured to select the LSP's path among equal-cost paths: Most-fill , Least-fill , or Random .	detail extensive
Signal type	Signal type for GMPLS LSPs. The signal type determines the peak data rate for the LSP: DS0 , DS3 , STS-1 , STM-1 , or STM-4 .	All levels
Encoding type	LSP encoding type: Packet , Ethernet , PDH , SDH/SONET , Lambda , or Fiber .	All levels
Switching type	Type of switching on the links needed for the LSP: Fiber , Lambda , Packet , TDM , or PSC-1 .	All levels
GPID	Generalized Payload Identifier (identifier of the payload carried by an LSP): HDLC , Ethernet , IPv4 , PPP , or Unknown .	All levels
Protection	Configured protection capability desired for the LSP: Extra , Enhanced , none , One plus one , One to one , or Shared .	All levels
Upstream label in	(Bidirectional LSPs) Incoming label for reverse direction traffic for this LSP.	All levels
Upstream label out	(Bidirectional LSPs) Outgoing label for reverse direction traffic for this LSP.	All levels

Table 37: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Suggested label received	(Bidirectional LSPs) Label the upstream node suggests to use in the Resv message that is sent.	All levels
Suggested label sent	(Bidirectional LSPs) Label the downstream node suggests to use in the Resv message that is returned.	All levels
Autobandwidth	(Ingress LSP) The LSP is performing autobandwidth allocation.	detail extensive
MinBW	(Ingress LSP) Configured minimum value of the LSP, in bps.	detail extensive
MaxBW	(Ingress LSP) Configured maximum value of the LSP, in bps.	detail extensive
Dynamic MinBW	(Ingress LSP) Displays the current dynamically specified minimum bandwidth allocation for the LSP, in bps.	detail extensive
Adjustment Timer	(Ingress LSP) Configured value for the adjust-timer statement, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	detail extensive
Adjustment Threshold	(Ingress LSP) Configured value for the adjust-threshold statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization.	detail extensive
Time for Next Adjustment	(Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.	detail extensive
Time of Last Adjustment	(Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.	detail extensive
Max AvgBW util	(Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps.	detail extensive
Overflow limit	(Ingress LSP) Configured value of the threshold overflow limit.	detail extensive
Overflow sample count	(Ingress LSP) Current value for the overflow sample count.	detail extensive
Bandwidth Adjustment in <i>nnn</i> second(s)	(Ingress LSP) Current value of the bandwidth adjustment timer, indicating the amount of time remaining until the bandwidth adjustment will take place, in seconds.	detail extensive
Underflow limit	(Ingress LSP) Configured value of the threshold underflow limit.	detail extensive
Underflow sample count	(Ingress LSP) Current value for the underflow sample count.	detail extensive
Underflow Max AvgBW	(Ingress LSP) The highest sample bandwidth among the underflow samples recorded currently. This is the signaling bandwidth if an adjustment occurs because of an underflow.	detail extensive

Table 37: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active path indicator	(Ingress LSP) A value of * indicates that the path is active. The absence of * indicates that the path is not active. In the following example, "long" is the active path. *Primary long Standby short	detail extensive
Primary	(Ingress LSP) Name of the primary path.	detail extensive
Secondary	(Ingress LSP) Name of the secondary path.	detail extensive
Standby	(Ingress LSP) Name of the path in standby mode.	detail extensive
State	(Ingress LSP) State of the path: Up or Dn (down).	detail extensive
COS	(Ingress LSP) Class-of-service value.	detail extensive
Bandwidth per class	(Ingress LSP) Active bandwidth for the LSP path for each MPLS class type, in bps.	detail extensive
Priorities	(Ingress LSP) Configured value of the setup priority and the hold priority respectively (the setup priority is displayed first), where 0 is the highest priority and 7 is the lowest priority. If you have not explicitly configured these values, the default values are displayed (7 for the setup priority and 0 for the hold priority).	detail extensive
OptimizeTimer	(Ingress LSP) Configured value of the optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	detail extensive
SmartOptimizeTimer	(Ingress LSP) Configured value of the smart optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	detail extensive
Reoptimization in xxx seconds	(Ingress LSP) Current value of the optimize timer, indicating the amount of time remaining until the path will be reoptimized, in seconds.	detail extensive
Computed ERO (S [L] denotes strict [loose] hops)	(Ingress LSP) Computed explicit route. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (S) or loose (L).	detail extensive
CSPF metric	(Ingress LSP) Constrained Shortest Path First metric for this path.	detail extensive

Table 37: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Received RRO	<p>(Ingress LSP) Received record route. A series of hops, each with an address followed by a flag. (In most cases, the received record route is the same as the computed explicit route. If Received RRO is different from Computed ERO, there is a topology change in the network, and the route is taking a detour.) The following flags identify the protection capability and status of the downstream node:</p> <ul style="list-style-type: none"> • 0x01—Local protection available. The link downstream from this node is protected by a local repair mechanism. This flag can be set only if the Local protection flag was set in the SESSION_ATTRIBUTE object of the corresponding Path message. • 0x02—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously). • 0x03—Combination of 0x01 and 0x02. • 0x04—Bandwidth protection. The downstream routing device has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section. • 0x08—Node protection. The downstream routing device has a backup path providing protection against link and node failure on the corresponding path section. If the downstream routing device can set up only a link-protection backup path, the Local protection available bit is set but the Node protection bit is cleared. • 0x09—Detour is established. Combination of 0x01 and 0x08. • 0x10—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engine LSP. This flag indicates to the ingress legacy edge router (LER) of this LSP that it should be rerouted. • 0x20—Node ID. Indicates that the address specified in the RRO's IPv4 or IPv6 sub-object is a node ID address, which refers to the router address or router ID. Nodes must use the same address consistently. • 0xb—Detour is in use. Combination of 0x01, 0x02, and 0x08. 	detail extensive
Index number	(Ingress LSP) Log entry number of each LSP path event. The numbers are in chronological descending order, with a maximum of 50 index numbers displayed.	extensive
Date	(Ingress LSP) Date of the LSP event.	extensive
Time	(Ingress LSP) Time of the LSP event.	extensive
Event	(Ingress LSP) Description of the LSP event.	extensive
Created	(Ingress LSP) Date and time the LSP was created.	extensive
Resv style	(Bypass) RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).	brief detail extensive
Labelin	Incoming label for this LSP.	brief detail
Labelout	Outgoing label for this LSP.	brief detail

Table 37: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
LSPname	Name of the LSP.	brief detail
Time left	Number of seconds remaining in the lifetime of the reservation.	detail
Since	Date and time when the RSVP session was initiated.	detail
Tspec	Sender's traffic specification, which describes the sender's traffic parameters.	detail
Port number	Protocol ID and sender or receiver port used in this RSVP session.	detail
PATH rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this router, and number of packets received from the upstream neighbor.	detail
PATH sentto	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor, and number of packets sent to the downstream routing device.	detail
RESV rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor. The output in this field, which is consistent with that in the PATH rcvfrom field, indicates that the RSVP negotiation is complete.	detail
Record route	Recorded route for the session, taken from the record route object.	detail
Soft preempt	Number of soft preemptions that occurred on a path and when the last soft preemption occurred. Only successful soft preemptions are counted (those that actually resulted in a new path being used).	detail
Soft preemption pending	Path is in the process of being soft preempted. This display is removed once the ingress router has calculated a new path.	detail
MPLS-TE LSP Defaults	Default settings for MPLS traffic engineered LSPs: <ul style="list-style-type: none"> • LSP Holding Priority—Determines the degree to which an LSP holds on to its session reservation after the LSP has been set up successfully. • LSP Setup Priority—Determines whether a new LSP that preempts an existing LSP can be established. • Hop Limit—Specifies the maximum number of routers the LSP can traverse (including the ingress and egress). • Bandwidth—Specifies the bandwidth in bits per second for the LSP. • LSP Retry Timer—Length of time in seconds that the ingress router waits between attempts to establish the primary path. 	defaults

The XML tag name of the **bandwidth** tag under the **auto-bandwidth** tag has been updated to **maximum-average-bandwidth**. You can see the new tag when you issue the **show mpls lsp extensive** command with the **| display xml** pipe option. If you have any scripts that use the **bandwidth** tag, ensure that they are updated to **maximum-average-bandwidth**.

Sample Output

show mpls lsp defaults

```
user@host> show mpls lsp defaults
MPLS-TE LSP Defaults
  LSP Holding Priority      0
  LSP Setup Priority       7
  Hop Limit                255
  Bandwidth                 0
  LSP Retry Timer          30 seconds
```

show mpls lsp descriptions

```
user@host> show mpls lsp descriptions
Ingress LSP: 3 sessions
To          LSP name          Description
10.0.0.195  to-sanjose                to-sanjose-desc
10.0.0.195  to-sanjose-other-desc      other-desc
Total 2 displayed, Up 2, Down 0
```

show mpls lsp detail

```
user@host> show mpls lsp detail
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
  10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
  20=Node-ID):
      10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 157, Since: Wed Jul 18 17:55:12 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 46128 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0
```



```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp detail (When Egress Protection Is in Standby Mode)

```
user@host> show mpls lsp detail
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
    10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
    20=Node-ID):
      10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 157, Since: Wed Jul 18 17:55:12 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 46128 protocol 0
Egress protection PLR as protector: Active

  PATH rcvfrom: 10.0.0.18 (1t-1/2/0.17) 3 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp detail (When Egress Protection Is in Effect During a Local Repair)

```
user@host> show mpls lsp detail
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
```

```

    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
        10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Down, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 157, Since: Wed Jul 18 17:55:12 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 46128 protocol 0
Egress protection PLR as protector: In Use
PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show mpls lsp extensive

```

user@host> show mpls lsp extensive
Ingress LSP: 4 sessions

1.1.1.1
  From: 3.3.3.3, State: Up, ActiveRoute: 0, LSPname: m120b-to-mx960
  ActivePath: DEFAULT (primary)
  FastReroute desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary  DEFAULT          State: Up
  Priorities: 7 0
  SmartOptimizeTimer: 180
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 310)
10.0.35.5 S 10.0.15.1 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
        10.0.34.4(flag=1) 10.0.14.1
50 Sep 13 16:08:19.712 Record Route: 10.0.35.5(flag=1) 10.0.15.1
49 Sep 13 16:08:16.720 Record Route: 10.0.34.4(flag=1) 10.0.14.1
48 Sep 13 16:08:16.699 Fast-reroute Detour Up
47 Sep 13 16:08:13.702 Record Route: 10.0.34.4 10.0.14.1
46 Sep 13 16:08:13.702 Up
45 Sep 13 16:08:13.672 Originate make-before-break call
44 Sep 13 16:08:13.672 CSPF: computation result accepted 10.0.34.4 10.0.14.1

43 Sep 13 16:08:13.672 Selected as active path
42 Sep 13 16:08:13.672 Make-before-break: Switched to new instance
41 Sep 13 16:08:01.685 Pending path switchover, skip CSPF run[3 times]
40 Sep 13 16:06:33.910 Deselected as active

```

```

39 Sep 13 16:06:33.910 Pending path switchover, skip CSPF run

38 Sep 13 16:06:19.521 Record Route: 10.0.35.5 10.0.15.1
37 Sep 13 16:06:19.518 ResvTear received
36 Sep 13 16:06:19.518 Fast-reroute Detour Down
35 Sep 13 16:06:16.676 Record Route: 10.0.35.5(flag=1) 10.0.15.1
34 Sep 13 16:06:13.670 Record Route: 10.0.35.5 10.0.15.1
33 Sep 13 16:06:13.670 Up
32 Sep 13 16:06:13.569 Pending path switchover, skip CSPF run

31 Sep 13 16:06:13.569 CSPF: link down/deleted:
10.0.34.3(3.3.3.3:79)(m120-b-re1.00/3.3.3.3)->0.0.0.0(0.0.0.0:0)(m120-b-re1.04/0.0.0.0)

30 Sep 13 16:06:13.552 Pending path switchover, skip CSPF run

29 Sep 13 16:06:13.552 CSPF: link down/deleted:
0.0.0.0(0.0.0.0:0)(m120-b-re1.04/0.0.0.0)->0.0.0.0(4.4.4.4:0)(m10i-a-re0.00/4.4.4.4)

28 Sep 13 16:06:13.549 Originate make-before-break call
27 Sep 13 16:06:13.549 CSPF: computation result accepted 10.0.35.5 10.0.15.1

26 Sep 13 16:06:13.548 Tunnel local repaired
25 Sep 13 16:06:13.546 Record Route: 10.0.23.2 10.0.12.1
24 Sep 13 16:06:13.546 10.0.34.3: Tunnel local repaired
23 Sep 13 16:06:13.546 10.0.34.3: Down
22 Sep 13 16:03:46.842 Fast-reroute Detour Up
21 Sep 13 16:03:42.730 Record Route: 10.0.34.4(flag=1) 10.0.14.1
20 Sep 13 16:03:39.836 Selected as active path
19 Sep 13 16:03:39.834 Record Route: 10.0.34.4 10.0.14.1
18 Sep 13 16:03:39.834 Up
17 Sep 13 16:03:39.698 Originate Call
16 Sep 13 16:03:39.698 CSPF: computation result accepted 10.0.34.4 10.0.14.1

15 Sep 13 16:03:39.697 Clear Call
14 Sep 13 16:03:39.696 Deselected as active
13 Sep 13 16:03:37.837 Record Route: 10.0.34.4 10.0.14.1
12 Sep 13 16:03:32.829 Fast-reroute Detour Down
11 Sep 13 16:02:15.493 Record Route: 10.0.34.4(flag=1) 10.0.14.1
10 Sep 13 16:02:15.486 Fast-reroute Detour Up
9 Sep 13 16:02:12.468 Record Route: 10.0.34.4 10.0.14.1
8 Sep 13 16:02:07.460 Fast-reroute Detour Down
7 Sep 13 15:57:46.741 Fast-reroute Detour Up
6 Sep 13 15:57:40.768 Record Route: 10.0.34.4(flag=1) 10.0.14.1
5 Sep 13 15:57:37.761 Selected as active path
4 Sep 13 15:57:37.760 Record Route: 10.0.34.4 10.0.14.1
3 Sep 13 15:57:37.760 Up
2 Sep 13 15:57:37.733 Originate Call
1 Sep 13 15:57:37.733 CSPF: computation result accepted 10.0.34.4 10.0.14.1

Created: Fri Sep 13 15:57:38 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 4 sessions, 6 detours
Total 0 displayed, Up 0, Down 0

Transit LSP: 6 sessions, 1 detours

1.1.1.1
From: 3.3.3.3, LSPstate: Up, ActiveRoute: 0
LSPname: m120b-to-mx960, LSPpath: Primary
Suggested label received: -, Suggested label sent: -

```

```

Recovery label received: -, Recovery label sent: 302288
Resv style: 1 FF, Label in: 300416, Label out: 302288
Time left: 147, Since: Fri Sep 13 16:08:16 2013
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 4 receiver 13955 protocol 0
Detour branch from 10.0.34.4, to skip 1.1.1.1, Up
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Adspec: received MTU 1500
  Path MTU: received 0
  PATH rcvfrom: 10.0.34.4 (ge-4/3/7.0) 7 pkts
  Adspec: received MTU 1500 sent MTU 1500
  PATH sentto: 10.0.35.5 (ge-3/1/0.0) 7 pkts
  RESV rcvfrom: 10.0.35.5 (ge-3/1/0.0) 7 pkts
  Explicit route: 10.0.35.5 10.0.15.1
  Record route: 10.0.34.3 10.0.34.4 <self>10.0.35.5 10.0.15.1
Label in: 300416, Label out: 302288
Total 1 displayed, Up 1, Down 0

```

show mpls lsp ingress extensive

```

user@host> show mpls lsp ingress extensive
Ingress LSP: 1 sessions

50.0.0.1
  From: 10.0.0.1, State: Up, ActiveRoute: 0, LSPName: test
  ActivePath: (primary)
  LSPType: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    OptimizeTimer: 300
    SmartOptimizeTimer: 180
    Reoptimization in 240 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
1.1.1.2 S 4.4.4.1 S 5.5.5.2 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    1.1.1.2 4.4.4.1 5.5.5.2
  17 Aug 3 13:17:33.601 CSPF: computation result ignored, new path less avail
bw[3 times]
  16 Aug 3 13:02:51.283 CSPF: computation result ignored, new path no benefit[2
times]
  15 Aug 3 12:54:36.678 Selected as active path
  14 Aug 3 12:54:36.676 Record Route: 1.1.1.2 4.4.4.1 5.5.5.2
  13 Aug 3 12:54:36.676 Up
  12 Aug 3 12:54:33.924 Deselected as active
  11 Aug 3 12:54:33.924 Originate Call
  10 Aug 3 12:54:33.923 Clear Call
  9 Aug 3 12:54:33.923 CSPF: computation result accepted 1.1.1.2 4.4.4.1
5.5.5.2
  8 Aug 3 12:54:33.922 2.2.2.2: No Route toward dest
  7 Aug 3 12:54:28.177 CSPF: computation result ignored, new path no benefit[4
times]
  6 Aug 3 12:35:03.830 Selected as active path
  5 Aug 3 12:35:03.828 Record Route: 2.2.2.2 3.3.3.2
  4 Aug 3 12:35:03.827 Up
  3 Aug 3 12:35:03.814 Originate Call
  2 Aug 3 12:35:03.814 CSPF: computation result accepted 2.2.2.2 3.3.3.2
  1 Aug 3 12:34:34.921 CSPF failed: no route toward 50.0.0.1

```

Created: Tue Aug 3 12:34:35 2010
Total 1 displayed, Up 1, Down 0

show mpls lsp extensive (automatic bandwidth adjustment enabled)

```

user@host> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPName: E-D
  ActivePath: (primary)
  Node/Link protection desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Autobandwidth
  MinBW: 300bps, MaxBW: 1000bps, Dynamic MinBW: 1000bps
  Adjustment Timer: 300 secs AdjustThreshold: 25%
  Max AvgBW util: 963.739bps, Bandwidth Adjustment in 0 second(s).
  Min BW Adjust Interval: 1000, MinBW Adjust Threshold (in %): 50
  Overflow limit: 0, Overflow sample count: 0
  Underflow limit: 0, Underflow sample count: 9, Underflow Max AvgBW: 614.421bps

  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    Bandwidth: 1000bps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
  10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
  20=Node-ID):
      192.168.0.6(flag=0x20) 10.0.0.18(Label=299792) 192.168.0.4(flag=0x20)
  10.0.0.22(Label=3)
    12 Apr 30 10:25:17.024 Make-before-break: Switched to new instance
    11 Apr 30 10:25:16.023 Record Route: 192.168.0.6(flag=0x20)
  10.0.0.18(Label=299792) 192.168.0.4(flag=0x20) 10.0.0.22(Label=3)
    10 Apr 30 10:25:16.023 Up
    9 Apr 30 10:25:16.023 Automatic Autobw adjustment succeeded: BW changes from
  300 bps to 1000 bps
    8 Apr 30 10:25:15.946 Originate make-before-break call
    7 Apr 30 10:25:15.946 CSPF: computation result accepted 10.0.0.18 10.0.0.22

    6 Apr 30 10:16:42.891 Selected as active path
    5 Apr 30 10:16:42.891 Record Route: 192.168.0.6(flag=0x20)
  10.0.0.18(Label=299776) 192.168.0.4(flag=0x20) 10.0.0.22(Label=3)
    4 Apr 30 10:16:42.890 Up
    3 Apr 30 10:16:42.828 Originate Call
    2 Apr 30 10:16:42.828 CSPF: computation result accepted 10.0.0.18 10.0.0.22

    1 Apr 30 10:16:14.064 CSPF: could not determine self[2 times]
  Created: Tue Apr 30 10:15:16 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show mpls lsp bypass extensive

```

user@host # show mpls lsp bypass extensive

Ingress LSP: 1 sessions

2.2.2.2
  From: 1.1.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->1.1.2.2
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 300032
  Resv style: 1 SE, Label in: -, Label out: 300032
  Time left: -, Since: Tue Dec 3 15:19:49 2013
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 55750 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 1
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 1.1.5.2 (lt-1/2/0.15) 1221 pkts
  RESV rcvfrom: 1.1.5.2 (lt-1/2/0.15) 1221 pkts, Entropy label: No
  Explct route: 1.1.5.2 1.2.5.1
  Record route: <self> 1.1.5.2 1.2.5.1
+   4 Dec 3 15:19:49 Record Route: 1.1.5.2 1.2.5.1
+   3 Dec 3 15:19:49 Up
+   2 Dec 3 15:19:49 CSPF: computation result accepted
+   1 Dec 3 15:19:47 Originate Call
Total 1 displayed, Up 1, Down 0
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
Transit LSP: 0 sessions

```

show mpls lsp p2mp

```

user@host> show mpls lsp p2mp
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1
To          From          State Rt P ActivePath      LSPname
10.255.245.51 10.255.245.50 Up    0 * path1        p2mp-branch-1
P2MP name: p2mp-lsp2, P2MP branch count: 1
To          From          State Rt P ActivePath      LSPname
10.255.245.51 10.255.245.50 Up    0 * path1        p2mp-st-br1
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show mpls lsp p2mp detail

```

user@host> show mpls lsp p2mp detail
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-branch-1

```

```

ActivePath: path1 (primary)
P2MP name: p2mp-lsp1
LoadBalance: Random
Encoding type: Packet, Switching type: Packet, GPID: IPv4
*Primary path1 State: Up
  Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
192.168.208.17 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
    192.168.208.17
P2MP name: p2mp-lsp2, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-st-br1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp2
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary path1 State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
192.168.208.17 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):
      192.168.208.17
Total 2 displayed, Up 2, Down 0

```

show mpls lsp detail count-active-routes

```

user@host> show mpls lsp detail count-active-routes
Ingress LSP: 1 sessions

213.119.192.2
  From: 156.154.162.128, State: Up, ActiveRoute: 1, LSPname: to-lahore
  ActivePath: (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Autobandwidth
  MinBW: 5Mbps MaxBW: 250Mbps
  Adjustment Timer: 300 secs
  Max AvgBW util: 60.2599Mbps, Bandwidth Adjustment in 0 second(s).
  Overflow limit: 0, Overflow sample count: 0
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Priorities: 7 0
    Bandwidth: 5Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 4)
10.252.0.177 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
      10.252.0.177
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show mpls lsp statistics extensive

```
user@host> show mpls lsp statistics extensive
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPName: E-D
  Statistics: Packets 302, Bytes 28992
  Aggregate statistics: Packets 302, Bytes 28992
  ActivePath: (primary)
  LSPType: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.0.0.18 10.0.0.22
    6 Oct  3 11:18:28.281 Selected as active path
    5 Oct  3 11:18:28.281 Record Route:  10.0.0.18 10.0.0.22
    4 Oct  3 11:18:28.280 Up
    3 Oct  3 11:18:27.995 Originate Call
    2 Oct  3 11:18:27.995 CSPF: computation result accepted  10.0.0.18 10.0.0.22

    1 Oct  3 11:17:59.118 CSPF failed: no route toward 192.168.0.4[2 times]
  Created: Wed Oct  3 11:17:01 2012
Total 1 displayed, Up 1, Down 0
```


show mpls lsp autobandwidth

Syntax	show mpls lsp autobandwidth <brief detail extensive> <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced in Junos OS Release 11.4. Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Display automatic bandwidth information for the LSP(s).
Options	<p>brief detail extensive — (Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.</p> <p>logical-system (all <i>logical-system-name</i>) — (Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show mpls lsp on page 1209
List of Sample Output	show mpls lsp autobandwidth on page 1228
Output Fields	Table 38 on page 1227 describes the output fields for the show mpls lsp autobandwidth command. Output fields are listed in the approximate order in which they appear.

Table 38: show mpls lsp autobandwidth Output Fields

Field Name	Field Description	Level of Output
To	Destination (egress routing device) of the session.	All Levels
From	Source (ingress routing device) of the session.	All Levels
LSPname	Name of the LSP.	All Levels
Min BW	(Ingress LSP) Configured minimum value of the LSP, in bps.	detail extensive
Max BW	(Ingress LSP) Configured maximum value of the LSP, in bps.	detail extensive
Max AvgBW util	(Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps.	detail extensive
Overflow limit	(Ingress LSP) Configured value of the threshold overflow limit.	detail extensive
Overflow sample count	(Ingress LSP) Current value for the overflow sample count.	detail extensive
Underflow limit	(Ingress LSP) Configured value of the threshold underflow limit.	detail extensive

Table 38: show mpls lsp autobandwidth Output Fields (*continued*)

Field Name	Field Description	Level of Output
Underflow sample count	(Ingress LSP) Current value for the underflow sample count.	detail extensive
Adjustment Timer	(Ingress LSP) Configured value for the adjust-timer statement, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	detail extensive
Adjustment Threshold	(Ingress LSP) Configured value for the adjust-threshold statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization.	detail extensive
Time for Next Adjustment	(Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.	detail extensive
Time of Last Adjustment	(Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.	detail extensive
Last BW	Previous active bandwidth of the LSP.	detail extensive
Last Requested BW	Bandwidth requested in the previous automatic bandwidth adjustment.	detail extensive
Last Signaled BW	Bandwidth signaled in the previous automatic bandwidth adjustment.	detail extensive
Highest Watermark BW	Maximum bandwidth used by the LSP.	detail extensive
Total AutoBw Adjustments	Total number of attempts to adjust automatic bandwidth including failed and successful adjustments.	detail extensive
Successful Adjustments	Number of successful automatic bandwidth adjustments.	detail extensive
Failed Adjustments	Number of failed automatic bandwidth adjustments.	detail extensive

Sample Output

show mpls lsp autobandwidth

```

user@host> show mpls lsp autobandwidth extensive
To: 10.255.106.133,
From: 10.255.106.135, LSPname: r0-r1
Min BW: 100kbps, Max BW: 0bps, Max AvgBW util: 2.33249Mbps
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 0
Adjustment Timer: 300 sec, Adjustment Threshold: 0
Time for Next Adjustment: 23 sec, Time of Last Adjustment: Fri Jun 3 21:05:37
2011
Last BW: 100kbps, Last Requested BW: 2.2169Mbps, Last Signaled BW: 2.2169Mbps,
Highest Watermark BW: 2.33249Mbps
Total AutoBw Adjustments: 1, Successful Adjustments: 1, Failed Adjustments: 0

```


show mpls path

List of Syntax [Syntax on page 1230](#)
[Syntax \(EX Series Switches\) on page 1230](#)

Syntax show mpls path
 <instance *instance-name*>
 <logical-system (all | *logical-system-name*)>
 <path-name>

Syntax (EX Series Switches) show mpls path
 <path-name>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.5 for EX Series switches.
instance *instance-name* option added in Junos OS Release 15.1.

Description Display dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options **none**—Display standard information about all MPLS LSPs.

instance *instance-name*—(Optional) Display the dynamic MPLS LSP for the specified instance. If *instance-name* is omitted, dynamic MPLS LSP for the master instance is displayed.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

path-name—(Optional) Display information about the specified LSP only.

Required Privilege Level view

List of Sample Output [show mpls path on page 1231](#)

Output Fields [Table 39 on page 1230](#) describes the output fields for the **show mpls path** command. Output fields are listed in the approximate order in which they appear.

Table 39: show mpls path Output Fields

Field Name	Field Description
Path name	Information about ingress LSPs. Each path has one line of output.
Address	Addresses of the routing devices that form the LSP.
Strict/loose address	Whether the address is configured as a strict or loose address.

Sample Output

show mpls path

```
user@host> show mpls path
Path name      Address          Strict/loose address
p1             123.456.55.6    Strict
               123.456.1.6     Loose
p2             191.456.1.4     Strict
```

show mpls srlg

Syntax `show mpls srlg`
`<logical-systems (all | logical-system-name)>`

Release Information Command introduced before Junos OS Release 11.4.

Description Display Shared Risk Link Group (SRLG) cost and value configuration information.



NOTE: If an SRLG is associated with a link that is used by an ingress LSP in the router, then on deleting the SRLG configuration from that router, the SRLG gets removed from the SRLG table only on the next reoptimization of the LSP. Until then, the output of the run `show mpls srlg` command displays `Unknown-XXX` instead of the SRLG name and a non zero `srlg-cost` for that SRLG.

Options `logical-system (all | logical-system-name)`—(Optional) View SRLG configuration information for all logical systems or a particular logical system.

Required Privilege Level view

Related Documentation

- [Example: Configuring SRLG on page 81](#)

Output Fields [Table 40 on page 1232](#) lists the output fields for the `show mpls srlg` command. Output fields are listed in the approximate order in which they appear.

Table 40: show mpls srlg Output Fields

Field Name	Field Description
SRLG	Name of the SRLG.
Value	A group ID for the SRLG ranging from 1 through 4294967295.
Cost	A cost for the Shared Risk Link Group (SRLG) ranging from 1 through 65535.

Sample Output

```
user@host> show mpls srlg
```

```
SRLG      Value      Cost
srlg-a    101        10
```

show mpls static-lsp

Syntax show mpls static-lsp
 <brief | detail | extensive | terse>
 <bypass>
 <descriptions>
 <down | up>
 <ingress>
 <instance *instance-name*>
 <logical-system (all | *logical-system-name*)>
 <lsp-type>
 <name *name*>
 <statistics>
 <transit>

Release Information Command introduced in Junos OS Release 10.1.
instance *instance-name* option added in Junos OS Release 15.1.
 Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
 Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.

Description Display information about configured and active static Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options **none**—Display standard information about all configured and active static MPLS LSPs.

brief | detail | extensive | terse—(Optional) Display the specified level of output. The **extensive** option displays the same information as the **detail** option, but covers the most recent 50 events.

bypass—(Optional) Display LSPs used for protecting other static LSPs.

descriptions—(Optional) Display the MPLS static LSP descriptions. To view this information, you must configure the description statement at the **[edit protocols mpls static-label-switched-path *path-name* bypass]**, **[edit protocols mpls static-label-switched-path *path-name* ingress]**, or **[edit protocols mpls static-label-switched-path *path-name* transit *incoming-label*]** hierarchy levels. Only static LSPs with a description are displayed.

down | up—(Optional) Display only static LSPs that are inactive or active, respectively.

instance *instance-name*—(Optional) Display information about all configured and active static MPLS LSPs for the specified routing instance. If ***instance-name*** is omitted, information about all configured and active static MPLS LSPs for the master instance is displayed.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-type—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.

- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

name *name*—(Optional) Display information about the specified static LSP or group of LSPs.

statistics—(Optional) Display accounting information about static LSPs.

transit—(Optional) Display static LSPs transiting this routing device.

Required Privilege Level view

List of Sample Output [show mpls static-lsp extensive on page 1235](#)
[show mpls static-lsp statistics ingress on page 1235](#)
[show mpls static-lsp \(when MPLS stitching is used\) on page 1235](#)

Output Fields [Table 30 on page 1191](#) describes the output fields for the **show mpls static-lsp** command. Output fields are listed in the approximate order in which they appear.

Table 41: show mpls static-lsp Output Fields

Field Name	Field Description	Level of Output
Ingress LSPs	Information about the static LSPs on the ingress routing device. Each session has one line of output.	All levels
Transit LSPs	Number of static LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
Bypass LSPs	Information about the bypass LSPs configured on the routing device. Each session has one line of output.	All levels
LSPname	Name of the static LSP.	All levels
To	Destination (egress routing device) of the session.	All levels
State	State of the static LSP handled by this RSVP session: Up , Dn (down), or Restart .	All levels
Packets	Number of packet transiting the static LSP (statistics option only).	All levels
Bytes	Number of bytes transiting the static LSP (statistics option only).	All levels
Nexthop	IP address for the next-hop router for the static LSP.	detail, extensive
Bypass	(Bypass LSP) Destination address (egress routing device) for the bypass LSP.	All levels
Link protection desired	Link protection has been requested by the ingress routing device.	detail, extensive
LabelOperation	Label operation to perform: Push , Pop , Swap .	detail, extensive

Table 41: show mpls static-lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Outgoing-label	Outgoing label to use for the MPLS packet in either push or swap label operations.	detail, extensive
Created	(Ingress LSP) Date and time the static LSP was created.	extensive
Bandwidth	Bandwidth configured for the static LSP.	detail, extensive
Resv style	(Bypass) RSVP reservation style. This field consists of two parts: the number of active reservations and the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).	All levels

Sample Output

show mpls static-lsp extensive

```

user@host> show mpls static-lsp extensive
Ingress LSPs:
LSPname: alpha-to-beta, To: 192.168.14.1
  State: Dn
  Nexthop: 192.168.10.1
  LabelOperation: Push, Outgoing-label: 1000001
  Created: Thu Jan 14 16:44:43 2010
  Bandwidth: 0 bps
Total 1, displayed 1, Up 0, Down 1

Transit LSPs:
Total 0, displayed 0, Up 0, Down 0

Bypass LSPs:
Total 0, displayed 0, Up 0, Down 0

```

show mpls static-lsp statistics ingress

```

user@host> show mpls static-lsp statistics ingress
Ingress LSPs:
LSPname           To           State    Packets    Bytes
alpha-to-beta     192.168.14.1 Dn        NA         NA
Total 1, displayed 1, Up 0, Down 1

```

show mpls static-lsp (when MPLS stitching is used)

The show mpls static-lsp command was extended in Junos release 14.1X53-D25 to accommodate the stitching feature of MPLS. This example shows the LSP state as 'InProgress' because the LSP is waiting for protocol next-hop resolution. For more information, see

```

user@host> show mpls static-lsp
Ingress LSPs:
Total 0, displayed 0, Up 0, Down 0
Transit LSPs: LSPname           Incoming-label  State
to-165        1000000        InProgress

```

show performance-monitoring mpls lsp


Syntax	show performance-monitoring mpls lsp <brief detail extensive> <name <i>lsp name</i> >
Release Information	Command introduced in Junos OS Release 15.1.
Description	Display the following performance monitoring data: <ul style="list-style-type: none"> • Packet loss measurement • Packet throughput measurement • Two-way channel delay • Round-trip delay • Inter-packet delay variation (IPDV)
Options	none —Display standard information performance monitoring data. brief detail extensive —(Optional) Display the specified level of output.
	<div>  <p>NOTE: The extensive option displays the same information as the detail option.</p> </div>
	name <i>lsp name</i> —(Optional) Display information about the specified LSP.
Required Privilege Level	View
Related Documentation	<ul style="list-style-type: none"> • clear performance-monitoring mpls lsp on page 1149 • performance-monitoring (Protocols MPLS) on page 922
List of Sample Output	show performance-monitoring mpls lsp on page 1238 show performance-monitoring mpls lsp detail on page 1239
Output Fields	Table 42 on page 1237 describes the output fields for the show performance-monitoring mpls lsp command. Output fields are listed in the approximate order in which they appear.

Table 42: show performance-monitoring mpls lsp Output Fields

Field Name	Display Data		Field Description	Level of Output
Session	Total		Total number of performance monitoring sessions created.	All Levels
	Up		Number of performance monitoring sessions that are up and running.	All Levels
	Down		Number of performance monitoring sessions that are down.	All Levels
LSP name			Name of the LSP.	All Levels
Loss measurement Data	Traffic-class		Traffic class for which loss measurement is performed.	All Levels
	Queries sent		Total number of queries sent for loss measurement.	All Levels
	Responses received		Total number of responses received for loss measurement queries.	All Levels
	Responses dropped due to errors		Total number of loss measurement responses dropped due to errors.	All Levels
	Queries timeout		Number of timed out queries sent for loss measurement.	All Levels
	Forward loss measurement	Average packet loss	Average packet loss (total loss of packets divided by the total number of samples used since the session is up).	All Levels
		Average packet throughput	Total number of packets sent divided by the time considered for measurement.	All Levels
	Reverse loss measurement	Average packet loss	Average packet loss (total loss of packets divided by the total number of samples used since the session is up).	All Levels
		Average packet throughput	Total number of packets sent divided by the time considered for measurement.	All Levels

Table 42: show performance-monitoring mpls lsp Output Fields (*continued*)

Field Name	Display Data	Field Description	Level of Output
Delay measurement Data	Traffic-class	Traffic class for which delay measurement is performed.	All Levels
	Queries sent	Total number of queries sent for delay measurement.	All Levels
	Responses received	Total number of responses received for delay measurement queries.	All Levels
	Responses dropped due to errors	Total number of delay measurement responses dropped due to errors.	All Levels
	Queries timeout	Number of timed out queries sent for delay measurement.	All Levels
	Best 2-way channel delay	Best available two-way channel delay.	All Levels
	Worst 2-way channel delay	Worst available two-way channel delay.	All Levels
	Best round trip time	Best available round-trip time.	All Levels
	Worst round trip time	Worst available round-trip time.	All Levels
	Avg absolute fw delay variation	Average of the variation in forward delay.	All Levels
	Avg absolute rv delay variation	Average of the variation in reverse delay.	All Levels
	Two-way channel delay	Sum of packet delays, excluding the processing time of the remote provider edge (PE) router.	detail, extensive
	Two-way round trip delay	Total time taken for completing round-trip of packet.	detail, extensive

Sample Output

show performance-monitoring mpls lsp

```

user@host> show performance-monitoring mpls lsp
Session Total: 3 Up: 3 Down: 0
LSP name:to_bad, PM State:Up
Loss measurement Data:
Duration: 00:04:43
Traffic-class: None
Queries sent: 282
Responses received: 282

```

```

Responses dropped due to errors: 0
Queries timeout: 0
Forward loss measurement:
  Average packet loss: 0
  Average packet throughput: 554338
Reverse loss measurement:
  Average packet loss: 0
  Average packet throughput: 1352077
LSP name:to_bad, PM State:Up
Delay measurement Data:
  Duration: 00:04:43
  Traffic-class: 0
  Queries sent: 282
  Responses received: 282
  Responses dropped due to errors: 0
  Queries timeout: 0
  Best 2-way channel delay: 72 usecs
  Worst 2-way channel delay: 365 usecs
  Best round trip time: 843 usecs
  Worst round trip time: 105523 usecs
  Avg absolute fw delay variation: 1619 usecs
  Avg absolute rv delay variation: 1619 usecs
LSP name:to_bad, PM State:Up
Loss measurement Data:
  Duration: 00:04:43
  Traffic-class: None
  Queries sent: 282
  Responses received: 282
  Responses dropped due to errors: 0
  Queries timeout: 0
  Forward loss measurement:
    Average packet loss: 0
    Average packet throughput: 553927
  Reverse loss measurement:
    Average packet loss: 0
    Average packet throughput: 1351531
Delay measurement Data:
  Best 2-way channel delay: 76 usecs
  Worst 2-way channel delay: 368 usecs
  Best round trip time: 1082 usecs
  Worst round trip time: 126146 usecs
  Avg absolute fw delay variation: 1618 usecs
  Avg absolute rv delay variation: 1619 usecs

```

show performance-monitoring mpls lsp detail

```

user@host> show performance-monitoring mpls lsp detail
Session Total: 3 Up: 3 Down: 0
LSP name:to_bad, PM State:Up
Loss measurement Data:
  Duration: 00:04:53
  Traffic-class: None
  Queries sent: 292
  Responses received: 292
  Responses dropped due to errors: 0
  Queries timeout: 0
  Forward loss measurement:
    Average packet loss: 0
    Average packet throughput: 554486
  Packet loss samples:
    00000000 00000000 00000000 00000000 00000000

```

```
Packet throughput samples:
00554002 00557550 00557717 00558822 00557107
Reverse loss measurement:
Average packet loss: 0
Average packet throughput: 1352406
Packet loss samples:
00000000 00000000 00000000 00000000 00000000
Packet throughput samples:
01351088 01365948 01353926 01362976 01358788
LSP name:to_bad, PM State:Up
Delay measurement Data:
Duration: 00:04:53
Traffic-class: 0
Queries sent: 292
Responses received: 292
Responses dropped due to errors: 0
Queries timeout: 0
Best 2-way channel delay: 72 usecs
Worst 2-way channel delay: 365 usecs
Best round trip time: 843 usecs
Worst round trip time: 105523 usecs
Avg absolute fw delay variation: 1683 usecs
Avg absolute rv delay variation: 1684 usecs
Two-way channel delay:
73 usecs 73 usecs 73 usecs 73 usecs 72 usecs
Two-way round trip delay:
922 usecs 2234 usecs 884 usecs 1121 usecs 1169 usecs
LSP name:to_bad, PM State:Up
Loss measurement Data:
Duration: 00:04:53
Traffic-class: None
Queries sent: 292
Responses received: 292
Responses dropped due to errors: 0
Queries timeout: 0
Forward loss measurement:
Average packet loss: 0
Average packet throughput: 554089
Packet loss samples:
00000000 00000000 00000000 00000000 00000000
Packet throughput samples:
00554007 00557548 00557713 00558547 00557385
Reverse loss measurement:
Average packet loss: 0
Average packet throughput: 1351914
Packet loss samples:
00000000 00000000 00000000 00000000 00000000
Packet throughput samples:
01358923 01352980 01362436 01223841 01496977
Delay measurement Data:
Best 2-way channel delay: 76 usecs
Worst 2-way channel delay: 368 usecs
Best round trip time: 1082 usecs
Worst round trip time: 126146 usecs
Avg absolute fw delay variation: 1682 usecs
Avg absolute rv delay variation: 1683 usecs
Two-way channel delay:
76 usecs 76 usecs 76 usecs 77 usecs 77 usecs
Two-way round trip delay:
107496 usecs 102369 usecs 104048 usecs 1433 usecs 103306 usecs
```


show ted database

List of Syntax	Syntax on page 1242 Syntax (EX Series Switches) on page 1242
Syntax	<pre>show ted database <brief detail extensive> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <<i>system-name</i>></pre>
Syntax (EX Series Switches)	<pre>show ted database <brief detail extensive> <<i>system-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>instance <i>instance-name</i> option added in Junos OS Release 15.1.</p>
Description	Display the entries in the Multiprotocol Label Switching (MPLS) traffic engineering database.
Options	<p>none—Display standard information about all entries in the traffic engineering database.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display routing instance information for the specified instance. If <i>instance-name</i> is omitted, information is displayed for the master instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>system-name</i>—(Optional) Display traffic engineering database information for a particular system.</p>
Required Privilege Level	view
List of Sample Output	show ted database brief on page 1245 show ted database detail on page 1245 show ted database extensive on page 1246
Output Fields	<p>Table 43 on page 1242 describes the output fields for the show ted database command. Output fields are listed in the approximate order in which they appear.</p>

Table 43: show ted database Output Fields

Field Name	Field Description	Level of Output
TED database	Number of nodes and pseudonodes participating in IS-IS and OSPF domain routing.	All levels

Table 43: show ted database Output Fields (*continued*)

Field Name	Field Description	Level of Output
ID	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. If the node contains a router ID, it is displayed in parentheses.	brief
NodeID	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	extensive
Type	Type of node. It can be either Rtr (router) or Net (pseudonode).	All levels
Age(s)	How long since the node was last refreshed, in seconds.	All levels
LnkIn	Number of nodes pointing toward this node.	All levels
LnkOut	Number of nodes to which this node points.	All levels
Protocol	Protocol that reported the node information: <ul style="list-style-type: none"> • IS-IS(1)—IS-IS Level 1. • IS-IS(2)—IS-IS Level 2. • OSPF (area-number)—OSPF from the specified area. 	All levels
To	Address on the far end of a link.	detail extensive
Local	Address of the local interface being used to reach the remote node.	detail extensive
Remote	Address of the interface on the remote node.	detail extensive
Local interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	detail extensive
Remote interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	detail extensive
Metric	Configured traffic engineering metric.	extensive
IGP metric	Configured interior gateway protocol metric.	extensive
Static BW	Total interface bandwidth in bps.	extensive
Reservable bandwidth	Subscription factor for the interface, which is the percentage of the link bandwidth that can be used for the RSVP reservation process. You configure this by including the subscription statement when configuring RSVP.	extensive
Available BW [priority]	(Must include diffserv-te statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each priority level. The bandwidth shown is for the entire interface, not for each individual LSP.	extensive

Table 43: show ted database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Diffserv-TE BW Model	Bandwidth constraint model used by the LSPs.	extensive
Available BW [TE-class]	(Must include the diffserv-te statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each traffic engineering class.	extensive
Static BW [CT-class]	Total interface bandwidth used by an MPLS traffic class, in bps.	extensive
Interface Switching Capability Descriptor (<i>n</i>)	<p>Information about the interface switching capability descriptor, which is a subtype length value (TLV) of the link TLV. <i>n</i> is the index number.</p> <ul style="list-style-type: none"> • Switching type—Type of switching to be performed on a particular link: <ul style="list-style-type: none"> • PSC-1—Packet switch-capable 1 • PSC-2—Packet switch-capable 2 • PSC-3—Packet switch-capable 3 • PSC-4—Packet switch-capable 4 • L2SC—Layer-2-switch-capable • TDM—Time-division-multiplexing-capable • LSC—Lambda switch-capable • FSC—Fiber switch-capable • Encoding type—Encoding of the LSP being requested: <ul style="list-style-type: none"> • Packet • Ethernet • ANSI/ETSI PDH • Reserved • SDH /SONET • Digital Wrapper • Lambda (photonic) • Fiber • FiberSDH/SONET • Maximum LSP BW [priority] bps—Maximum LSP bandwidth information. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. <ul style="list-style-type: none"> • [<i>n</i>]—Priority level. The range is from 0 (high) through 7 (low). • <i>n</i> Mbps—Amount of the maximum bandwidth. • Minimum LSP BW—Minimum LSP bandwidth in Mbps. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. Minimum LSP BW is displayed only when switching type is PSC-1 or TDM. • Interface MTU—Displayed only when switching type is TDM. • Interface supports standard SONET/SDH—Displayed only when switching type is TDM. 	extensive

Sample Output

show ted database brief

```

user@host> show ted database brief
TED database: 12 ISIS nodes 0 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-A.00                      ---  3178    2      0
Router-B.00                      ---  3152    2      0
Router-B.02                      Net   802    0      2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-C.00                      ---  3126    2      0
Router-C.02                      Net   38     0      2 IS-IS(2)
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-D.00                      ---  3144    2      0
Router-D.02                      Net   723    0      2 IS-IS(2)
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-D.03                      Net   607    0      2 IS-IS(2)
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-E.00                      ---  3178    2      0
Router-E.02                      Net   131    0      2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-F.00                      ---  3153    2      0
Router-F.02                      Net   769    0      2 IS-IS(2)
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0

```

show ted database detail

```

TED database: 12 ISIS nodes 0 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-A.00                      ---  2913    2      0
Router-B.00                      ---  2887    2      0
Router-B.02                      Net   537    0      2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol

```

```

Router-C.00          ---      2861      2      0
Router-C.02          Net       597      0      2 IS-IS(2)
  To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                  Type Age(s) LnkIn LnkOut Protocol
Router-D.00          ---      2879      2      0
Router-D.02          Net       458      0      2 IS-IS(2)
  To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                  Type Age(s) LnkIn LnkOut Protocol
Router-D.03          Net       342      0      2 IS-IS(2)
  To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                  Type Age(s) LnkIn LnkOut Protocol
Router-E.00          ---      2913      2      0
Router-E.02          Net       640      0      2 IS-IS(2)
  To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                  Type Age(s) LnkIn LnkOut Protocol
Router-F.00          ---      2888      2      0
Router-F.02          Net       504      0      2 IS-IS(2)
  To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0

```

show ted database extensive

```

user@host> show ted database extensive
TED database: 12 ISIS nodes 0 INET nodes
NodeID: Router-A.00
  Type: ---, Age: 3067 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-B.00
  Type: ---, Age: 3041 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-B.02
  Type: Net, Age: 691 secs, LinkIn: 0, LinkOut: 2
  Protocol: IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
      Metric: 0
      IGP metric: 10
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
          [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
      Metric: 0
      IGP metric: 20
      Interface Switching Capability Descriptor(1):
        Switching type: Packet

```

```

        Encoding type: Packet
        Maximum LSP BW [priority] bps:
            [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
            [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
NodeID: Router-C.00
Type: ---, Age: 3015 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-C.02
Type: Net, Age: 751 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10      Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
NodeID: Router-D.00
Type: ---, Age: 3034 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-D.02
Type: Net, Age: 613 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
NodeID: Router-D.03
Type: Net, Age: 497 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):

```

```

Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
    [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
    [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
        [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
NodeID: Router-E.00
Type: ---, Age: 3068 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-E.02
Type: Net, Age: 21 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
        [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
        [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
NodeID: Router-F.00
Type: ---, Age: 3043 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-F.02
Type: Net, Age: 659 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
        [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:

```

[0] 0bps	[1] 0bps	[2] 0bps	[3] 0bps
[4] 0bps	[5] 0bps	[6] 0bps	[7] 0bps

show ted link

List of Syntax	Syntax on page 1250 Syntax (EX Series Switches) on page 1250
Syntax	<pre>show ted link <brief detail> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	<pre>show ted link <brief detail></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>instance <i>instance-name</i> option added in Junos OS Release 15.1.</p>
Description	Display Multiprotocol Label Switching (MPLS) traffic engineering database link information.
Options	<p>none—Display standard information about traffic engineering database link information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display routing instance information for the specified instance. If <i>instance-name</i> is omitted, information is displayed for the master instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show ted link brief on page 1251 show ted link detail on page 1251
Output Fields	Table 44 on page 1250 describes the output fields for the show ted link command. Output fields are listed in the approximate order in which they appear.

Table 44: show ted link Output Fields

Field Name	Field Description	Level of Output
ID	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	brief
-->ID	Hostname and address of the node that the link is going to. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	brief

Table 44: show ted link Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>hostname</i>	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	detail
<i>hostname</i>	Hostname and address of the node that the link is going to. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	detail
Local Path	Number of paths CSPF on the local routing device has placed on the link.	All levels
Metric	Configured traffic engineering metric.	extensive
IGP metric	Configured interior gateway protocol metric.	detail
Local BW	Amount of bandwidth the local routing device has placed on the link.	All levels
Local	Address of the local interface being used to reach the remote node.	detail extensive
Remote	Address of the interface on the remote node.	detail extensive
Local interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	detail
Remote interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	detail

Sample Output

show ted link brief

```

user@host> show ted link brief
ID                               ->ID                               LocalPath LocalBW
Router-B.02                     Router-A.00                       0 0bps
Router-B.02                     Router-B.00                       0 0bps
Router-C.02                     Router-B.00                       0 0bps
Router-C.02                     Router-C.00                       0 0bps
Router-D.02                     Router-F.00                       0 0bps
Router-D.02                     Router-D.00                       0 0bps
Router-D.03                     Router-D.00                       0 0bps
Router-D.03                     Router-C.00                       0 0bps
Router-E.02                     Router-A.00                       0 0bps
Router-E.02                     Router-E.00                       0 0bps
Router-F.02                     Router-E.00                       0 0bps
Router-F.02                     Router-F.00                       0 0bps

```

show ted link detail

```

user@host> show ted link detail
Router-B.02->Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps

```

```
    localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-B.02->Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 20 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-C.02->Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 40 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-C.02->Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-D.02->Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-D.02->Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 60 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-D.03->Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-D.03->Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-E.02->Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 60 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-E.02->Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 20 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-F.02->Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-F.02->Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 40 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
```

show ted protocol

List of Syntax	Syntax on page 1253 Syntax (EX Series Switches) on page 1253
Syntax	<pre>show ted protocol <brief detail> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	<pre>show ted protocol <brief detail></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>instance <i>instance-name</i> option added in Junos OS Release 15.1.</p>
Description	Display information about the protocols from which the Multiprotocol Label Switching (MPLS) traffic engineering database learned about its nodes.
Options	<p>none—Display standard information about the protocols from which the traffic engineering database learned about its nodes.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display routing instance information for the specified instance. If <i>instance-name</i> is omitted, information is displayed for the master instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show ted protocol on page 1254
Output Fields	<p>Table 45 on page 1253 describes the output fields for the show ted protocol command. Output fields are listed in the approximate order in which they appear.</p>

Table 45: show ted protocol Output Fields

Field Name	Field Description
Protocol name	<p>Protocol that reported the node information:</p> <ul style="list-style-type: none"> IS-IS(1)—IS-IS Level 1. IS-IS(2)—IS-IS Level 2. OSPF (<i>area-number</i>)—OSPF from the specified area.
Credibility	If the protocols provide conflicting information about a node, the protocol with the highest credibility value is the one that the traffic engineering database uses.
Self node	Address the protocol uses as the local address.

Sample Output

show ted protocol

```
user@host> show ted protocol
Protocol name      Credibility Self node
IS-IS(2)           2 (highest) corriedale.00(123.456.1.11)
IS-IS(1)           1           corriedale.00(123.456.1.11)
```

traceroute mpls bgp

Syntax `traceroute mpls bgp fec`
`<destination destination-address>`
`<detail>`
`<exp exp>`
`<fanout fanout-number>`
`<logical-system logical system-name>`
`<no-resolve>`
`<paths paths-number>`
`<pipe-mode>`
`<retries retries-number>`
`<routing-instance routing-instance-name>`
`<source source-address>`
`<ttl value>`
`<wait seconds>`

Release Information Command introduced in Junos OS Release 14.2.

Description Trace route to a remote host for an MPLS label-switched path (LSP) signaled by the Border Gateway Protocol (BGP). Use **traceroute mpls bgp** as a debugging tool to locate MPLS BGP forwarding issues in a network. (Currently supported for IPv4 packets only.)



NOTE: The `traceroute mpls bgp fec` command only supports single paths.

Options **fec**—Specify the IP address and optional prefix of the forwarding equivalence class (FEC). Suppose you are at PE1, you would want to use the IP address of PE2 to trace the BGP path to that router.

destination destination-address—(Optional) Specify the destination address to use when sending probes.

detail—(Optional) Display detailed output.

exp exp—(Optional) Specify the class of service to use when sending probes.

Range: 0 through 7

Default: 7

fanout fanout-number—(Optional) Specify the maximum number of next hops to search per node.

Range: 1 through 16

Default: 16

logical-system logical-system-name—(Optional) Specify the name of the logical system for the traceroute attempt.

no-resolve—(Optional) Specify not to resolve the hostname that corresponds to the IP address.

paths *paths-number*—(Optional) Specify the number of paths to search.

Range: 1 through 255

Default: 16

pipe-mode—(Optional) Specify to trace only the outermost FEC.

retries *retries-number*—(Optional) Specify the number of times to resend probe values.

Range: 1 through 9

Default: 3

routing-instance *routing-instance-name*—(Optional) Specify the name of the routing instance for the trace route attempt.

source *source-address*—(Optional) Specify the source address of the outgoing traceroute packets.

ttl *value*—(Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds.

Range: 1 through 125

Default: 64

wait *seconds*—(Optional) Specify the number of seconds to wait before resending a probe.

Range: 5 through 15

Default: 10

Required Privilege Level network

Related Documentation • [ping mpls bgp on page 1163](#)

List of Sample Output [traceroute mpls bgp on page 1257](#)
[traceroute mpls bgp detail on page 1257](#)

Output Fields [Table 46 on page 1256](#) describes the output fields for the **traceroute mpls bgp fec** command and the **traceroute mpls bgp fec detail** command. Output fields are listed in the approximate order in which they appear.

Table 46: traceroute mpls bgp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the traceroute mpls bgp fec command.	All levels
ttl	Time to live value of the labeled packet.	None
Label	Outgoing label used for forwarding the packet along the label-switched paths.	None

Table 46: traceroute mpls bgp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Protocol	Signaling protocol used. For this command, it is BGP.	None
Address	Address of the next hop.	None
Previous Hop	Address of the previous hop. Previous hop address of the first hop is null .	None
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	None
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	detail
Parent	Address of the previous hop. Parent value for the first hop is null .	detail
Return Code	Return code for reporting the result of processing the echo request by the receiver.	detail
Response time	Time for the echo request to reach the receiver.	detail
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is none .	detail
Label Stack	Label stack used to forward the packet.	detail

Sample Output

traceroute mpls bgp

```

user@host> traceroute mpls bgp fec
Probe options: retries 3, exp 7
ttl Label Protocol Address Previous Hop Probe Status Fec-Stack-Sent Fec-Change
1 299824 LDP 81.1.2.2 (null) Success LDP, BGP PUSH-RSVP
2 299825 RSVP 81.2.3.3 81.1.2.2 Success RSVP, LDP, BGP (null)
3 299826 RSVP 81.3.4.4 81.2.3.3 Egress RSVP, LDP, BGP POP-RSVP
3 299826 LDP 81.3.4.4 81.2.3.3 Success LDP, BGP (null)
4 299827 LDP 81.4.5.5 81.3.4.4 Egress LDP, BGP POP-LDP
4 299827 BGP 81.4.5.5 81.3.4.4 Egress BGP (null)

```

traceroute mpls bgp detail

```

user@host> traceroute mpls bgp fec detail
Probe options: retries 3, exp 7
Hop 2.2.1.81.rev.sfr.net (81.1.2.2) Depth 1
Probe status: Success
Parent: (null)
Return code: Label switched at stack-depth 1
Sender timestamp: 2013-03-22 05:55:19 PDT 822.99 msec

```

Receiver timestamp: 2013-03-22 05:55:19 PDT 856.05 msec
Response time: 33.06 msec
MTU: Unknown
Multipath type: IP bitmask
Address Range 1: 127.0.0.64 ~ 127.0.0.127
Label Stack:
Label 1 Value 299824 Protocol LDP
Label 2 Value 299276 Protocol BGP
Fec-Stack-Sent: LDP, BGP
Fec-Change:
Operation: PUSH Protocol RSVP

CHAPTER 30

RSVP Operational Commands

- `clear rsvp session`
- `clear rsvp statistics`
- `monitor label-switched-path`
- `ping mpls rsvp`
- `show rsvp interface`
- `show rsvp neighbor`
- `show rsvp session`
- `show rsvp statistics`
- `show rsvp version`
- `traceroute mpls rsvp`

clear rsvp session

List of Syntax [Syntax on page 1260](#)
 [Syntax \(EX and QFX Series Switches\) on page 1260](#)

Syntax clear rsvp session
 <connection-destination *address*>
 <connection-source *address*>
 <gracefully>
 <logical-system (all | *logical-system-name*)>
 <lsp-id *identifier*>
 <name *name*>
 <optimize-fast-reroute>
 <tunnel-id *identifier*>

Syntax (EX and QFX Series Switches) clear rsvp session
 <connection-destination *address*>
 <connection-source *address*>
 <gracefully>
 <lsp-id *identifier*>
 <name *name*>
 <optimize-fast-reroute>
 <tunnel-id *identifier*>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.5 for EX Series switches.
 Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description Reset and restart Resource Reservation Protocol (RSVP) sessions.

Options **none**—Reset and restart all RSVP sessions for which this routing device is the ingress, transit, or egress routing device.

connection-source *address*—(Optional) Source address for GMPLS and MPLS LSPs from the RSVP sender template.

connection-destination *address*—(Optional) Destination address for GMPLS and MPLS LSPs from the RSVP sender template.

gracefully—(Optional) Gracefully reset an RSVP session for a nonpacket LSP in two passes. In the first pass, the Admin-Status object is signaled along the path to the other endpoint of the RSVP session. In the second pass, the path used by the RSVP session is torn down. This option can only be used on the ingress or egress routing device of the RSVP session and is only valid for nonpacket LSPs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-id *identifier*—(Optional) LSP identifier (source port) for the RSVP sender template.

name *name*—(Optional) Reset and restart the specified RSVP session.

optimize-fast-reroute—(Optional) Begin fast reroute optimization.

tunnel-id *identifier*—(Optional) Tunnel identifier (destination port) for the RSVP session.

Required Privilege Level clear

Related Documentation

- [clear mpls lsp on page 1145](#)
- [show rsvp session on page 1281](#)

List of Sample Output [clear rsvp session on page 1261](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

[clear rsvp session](#)

```
user@host> clear rsvp session
```

clear rsvp statistics

List of Syntax	Syntax on page 1262 Syntax (EX Series Switches) on page 1262
Syntax	clear rsvp statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	clear rsvp statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Clear Resource Reservation Protocol (RSVP) packet and error statistics.
Options	none —Clear RSVP packet and error statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show rsvp statistics on page 1291
List of Sample Output	clear rsvp statistics on page 1262
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear rsvp statistics

```
user@host> clear rsvp statistics
```

monitor label-switched-path

Syntax `monitor label-switched-path lsp-name`
`<logical-system (logical-system-name)>`

Release Information Command introduced before Junos OS Release 7.4.
 Logical system support introduced in Junos OS Release 9.4.
 Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

Description Display the real-time status of the specified RSVP label-switched path (LSP). You can also use this command to monitor LSPs configured within logical systems.

Options `logical-system (logical-system-name)`—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-name—Name of the LSP.

Additional Information You can track the amount of traffic traversing an RSVP LSP and observe its essential parameters, such as uptime, ingress and egress addresses, labels, routes, and ports. Values are typically sampled every second. The display also allows you to scroll to other currently running LSPs. You cannot use this command to display information about static LSPs or LDP-signaled LSPs.

The output of this command shows how much each field has changed since you started the command or since you cleared the counters by using the `c` key. To control the output of the **monitor label-switched-path** command while it is running, use the keys listed in [Table 47 on page 1263](#). The keys are not case-sensitive.

Table 47: Output Control Keys for the monitor label-switched-path Command

Key	Action
c	Clears the screen and refreshes the display for this LSP.
f	Freezes the display, preventing new information from being displayed.
l	Monitors a different LSP. After you type l, you can type the new LSP name.
n	Displays information about the next LSP (whose name is alphabetically higher than the current LSP name) configured on the router.
p	Goes to the previous LSP (whose name is alphabetically lower than the current LSP name) configured on the router.
q or Esc	Quits the command and returns to the command prompt.
t	Thaws, or restarts, the data display for this LSP.

Required Privilege Level trace

List of Sample Output [monitor label-switched-path on page 1265](#)

Output Fields [Table 48 on page 1264](#) describes the output fields for the **monitor label-switched-path** command. Output fields are listed in the approximate order in which they appear.

Table 48: monitor label-switched-path Output Fields

Field Name	Field Description
(1)	Displays the following information: <ul style="list-style-type: none"> • hostname—Name of the router. • Seconds—Time elapsed since this display was started. • Time—Current local time.
(2)	Delay —Length of the time delay, in milliseconds, required to obtain the information in the monitor display. The first number shows the current sampling delay. The second number shows the shortest delay recorded to date. The third number shows the worst delay recorded to date. This delay can vary substantially depending on the system load.
(3)	Displays the following: <ul style="list-style-type: none"> • To—Destination address of the LSP. • From—Originating address of the LSP. • State—Current state of the LSP: Up or Down.
(4)	Displays the following: <ul style="list-style-type: none"> • LSPName—Name of the LSP. • Type—Type of LSP: Ingress, Egress, or Transit.
(5)	Displays the following: <ul style="list-style-type: none"> • Label in—Incoming label of the LSP. • Label out—Outgoing label of the LSP.
(6)	Port number —Port number for the sending router, the port number for the receiving router, and the protocol ID. For MPLS traffic engineering applications, the protocol ID is always 0.
(7/8)	Record route —All intermediate and egress router addresses for this LSP.
(9/10/11)	Displays traffic statistics: <ul style="list-style-type: none"> • Output packets—Number of packets that have traversed this LSP, and the change (delta) in the number since the last sample, typically 1 second ago. • Output bytes—Number of bytes that have traversed this LSP, and the change (delta) in the number since the last sample, typically 1 second ago.
(12)	Displays any errors the router encountered while attempting to retrieve information on the LSP.
(13)	Lists the keyboard commands you can use to navigate to other LSPs. For a description of the keyboard commands, see Table 47 on page 1263 .

Sample Output

monitor label-switched-path

```
user@host> monitor label-switched-path
(1) host                               Seconds: 112           Time: 15:32:22
(2)                                     Delay: 0/0/0
(3) To 10.10.10.16, From 10.10.10.17, state: Up
(4)  LSPname: k, type: Ingress
(5)  Label in: -, Label out: 126000
(6)  Port number: sender 1, receiver 45583, protocol 0
(7)  Record Route: <self> 192.168.224.196
(8)    192.168.224.202 192.168.224.179
(9)  Traffic statistics:
(10)    Output packets:                0                      [0]
(11)    Output bytes:                  0                      [0]
(12)
(13)Next='n', Prev='p', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c',
    LSP='l'
```

ping mpls rsvp

Syntax ping mpls rsvp
 <lsp-name>
 <count count>
 <destination address>
 <detail>
 <dynamic-bypass>
 <egress egress-address>
 <exp forwarding-class>
 <interface interface-name>
 <logical-system (all | logical-system-name)>
 <manual-bypass>
 <multipoint>
 <size bytes>
 <source source-address>
 <standby standby-path-name>
 <sweep>

Release Information Command introduced before Junos OS Release 7.4.
 The **egress** and **multipoint** options were introduced in Junos OS Release 9.2.
 The **size** and **sweep** options were introduced in Junos OS Release 9.6.
 The **dynamic-bypass** and **manual-bypass** options were introduced in Junos OS Release 10.2.

Description Check the operability of MPLS RSVP-signaled label-switched path (LSP) connections. Type Ctrl+c to interrupt a **ping mpls** command.

Options **count count**—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

destination address—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.



NOTE: When using the **detail** option, the reported time is based on the system time configured on the local and remote routers. Differences in these system times can result in inaccurate one way ping trip times being reported.

In practice, it is difficult to synchronize the system times of independent Juniper Networks routers with sufficient accuracy to provide a meaningful time value for the **detail** option (even when synchronized using NTP).

dynamic-bypass—(Optional) Ping dynamically generated bypass LSPs, used for protecting other LSPs.

egress *egress-address*—(Optional) Only the specified egress router or switch responds to the ping request.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

interface—(Optional) Specify the name of the interface protected by the manual bypass LSP. This option is only available when you have also used the **manual-bypass** option.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

lsp-name—Ping an RSVP-signaled LSP using an LSP name.

manual-bypass—(Optional) Ping manually configured bypass LSPs, used for protecting other LSPs. For this option, you must also specify the interface protected by the manual bypass LSP using the **interface** option.

multipoint—(Optional) Send ping requests to each of the egress routers or switches participating in a point-to-multipoint LSP. You can also include the **egress** option to ping a specific egress router or switch participating in a point-to-multipoint LSP.

size *bytes*—(Optional) Size of the LSP ping request packet (100 through 65468 bytes). Packets are 4-byte aligned. For example, if you enter a size of 101, 102, 103, or 104, the router or switch uses a size value of 104 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 100-byte minimum.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface.

standby *standby-path-name*—(Optional) Name of the standby path.

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level network

List of Sample Output [ping mpls rsvp \(Echo Reply Received\) on page 1268](#)
[ping mpls rsvp \(Echo Reply with Error Code\) on page 1268](#)

[ping mpls rsvp detail on page 1268](#)

[ping mpls rsvp multipoint egress detail count on page 1268](#)

[ping mpls rsvp multipoint detail count on page 1268](#)

[ping mpls rsvp destination detail count size on page 1269](#)

[ping mpls rsvp destination detail sweep size on page 1269](#)

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls rsvp (Echo Reply Received)

```
user@host> ping mpls rsvp test1
!!!!!--- lsping statistics ---5 packets transmitted, 5 packets received, 0% packet
loss
```

ping mpls rsvp (Echo Reply with Error Code)

```
user@host> ping mpls rsvp test2
!!xxx--- lsping statistics ---5 packets transmitted, 2 packets received, 60%
packet loss3 packets received with error status, not counted as received.
```

ping mpls rsvp detail

```
user@host> ping mpls rsvp to-green detail
Request for seq 1, to interface 67, labels <100095, 0, 0>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, labels <100095, 0, 0>
Reply for seq 2, return code: Egress-ok
```

ping mpls rsvp multipoint egress detail count

```
user@host>ping mpls rsvp sample-lsp multipoint egress 192.168.1.3 detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.242 ms
Local transmit time: 1205310695s 215737us
Remote receive time: 1205310695s 215979us

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

ping mpls rsvp multipoint detail count

```
user@host>ping mpls rsvp sample-lsp multipoint detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, return code: Unknown TLV, time: 9.877 ms
Local transmit time: 1205310615s 347317us
Remote receive time: 1205310615s 357194us
Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.351 ms
```

```

Local transmit time: 1205310615s 347262us
Remote receive time: 1205310615s 347613us
Reply for seq 1, egress 192.168.1.13, return code: Egress-ok, time: 0.301 ms
Local transmit time: 1205310615s 347167us
Remote receive time: 1205310615s 347468us
Timeout for seq 1, egress 192.168.1.1
Timeout for seq 1, egress 192.168.1.4
Timeout for seq 1, egress 192.168.1.14

--- lsping, egress 192.168.1.1 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.4 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.13 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.14 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

```

ping mpls rsvp destination detail count size

```

user@host> ping mpls rsvp chaser-access destination 192.168.0.1 detail count 1 size 4468

Request for seq 1, to interface 88, label 299984, packet size 4468
Reply for seq 1, return code: Egress-ok, time: 44.804 ms
    Local transmit time: 2009-03-30 22:05:02 CEST 408.629 ms
    Remote receive time: 2009-03-30 22:05:02 CEST 453.433 ms

--- lsping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

```

ping mpls rsvp destination detail sweep size

```

user@router> ping mpls rsvp chaser-access destination 192.168.0.1 detail sweep size 4500
Request for seq 1, to interface 86, no label stack., packet size 100
Reply for seq 1, return code: Egress-ok, time: -39.264 ms
    Local transmit time: 2009-04-24 14:05:40 CEST 541.423 ms
    Remote receive time: 2009-04-24 14:05:40 CEST 502.159 ms
Request for seq 2, to interface 86, no label stack., packet size 2300
Reply for seq 2, return code: Egress-ok, time: -38.179 ms
    Local transmit time: 2009-04-24 14:05:41 CEST 544.240 ms
    Remote receive time: 2009-04-24 14:05:41 CEST 506.061 ms
Request for seq 3, to interface 86, no label stack., packet size 4500
Timeout for seq 3
Request for seq 4, to interface 86, no label stack., packet size 3400
Reply for seq 4, return code: Egress-ok, time: -37.545 ms
    Local transmit time: 2009-04-24 14:05:45 CEST 549.953 ms
    Remote receive time: 2009-04-24 14:05:45 CEST 512.408 ms
Request for seq 5, to interface 86, no label stack., packet size 3952
Reply for seq 5, return code: Egress-ok, time: -37.176 ms
    Local transmit time: 2009-04-24 14:05:46 CEST 555.881 ms
    Remote receive time: 2009-04-24 14:05:46 CEST 518.705 ms
Request for seq 6, to interface 86, no label stack., packet size 4228
Reply for seq 6, return code: Egress-ok, time: -36.962 ms
    Local transmit time: 2009-04-24 14:05:47 CEST 561.809 ms
    Remote receive time: 2009-04-24 14:05:47 CEST 524.847 ms

```

```
Request for seq 7, to interface 86, no label stack., packet size 4368
Reply for seq 7, return code: Egress-ok, time: -36.922 ms
    Local transmit time: 2009-04-24 14:05:48 CEST 568.738 ms
    Remote receive time: 2009-04-24 14:05:48 CEST 531.816 ms
Request for seq 8, to interface 86, no label stack., packet size 4440
Reply for seq 8, return code: Egress-ok, time: -36.855 ms
    Local transmit time: 2009-04-24 14:05:49 CEST 575.669 ms
    Remote receive time: 2009-04-24 14:05:49 CEST 538.814 ms
Request for seq 9, to interface 86, no label stack., packet size 4476
Timeout for seq 9
Request for seq 10, to interface 86, no label stack., packet size 4460
Reply for seq 10, return code: Egress-ok, time: -36.906 ms
    Local transmit time: 2009-04-24 14:05:53 CEST 584.382 ms
    Remote receive time: 2009-04-24 14:05:53 CEST 547.476 ms
Request for seq 11, to interface 86, no label stack., packet size 4480
Timeout for seq 11
Request for seq 12, to interface 86, no label stack., packet size 4472
Timeout for seq 12
Request for seq 13, to interface 86, no label stack., packet size 4468
Reply for seq 13, return code: Egress-ok, time: -36.943 ms
    Local transmit time: 2009-04-24 14:06:00 CEST 594.884 ms
    Remote receive time: 2009-04-24 14:06:00 CEST 557.941 ms
Request for seq 14, to interface 86, no label stack., packet size 4476
Timeout for seq 14
Request for seq 15, to interface 86, no label stack., packet size 4472
Timeout for seq 15

--- lsp ping sweep result---
Maximum Transmission Unit (MTU) is 4468 bytes
```

show rsvp interface

List of Syntax	Syntax on page 1271 Syntax (EX Series Switches) on page 1271
Syntax	<pre>show rsvp interface <brief detail extensive> <instance <i>instance-name</i>> <link-management> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	<pre>show rsvp interface <brief detail extensive> <link-management></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>instance <i>instance-name</i> option added in Junos OS Release 15.1.</p>
Description	Display the status of Resource Reservation Protocol (RSVP)-enabled interfaces and packet statistics.
Options	<p>none—Display standard information about the status of RSVP-enabled interfaces and packet statistics.</p> <p>brief detail extensive link-management—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display RSVP status information for the specified instance. If <i>instance-name</i> is omitted, RSVP status information is displayed for the master instance.</p> <p>link-management—(Optional) Use the link-management option to display the control peers and corresponding TE-link information created by the Link Management Protocol (LMP).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show rsvp interface brief on page 1274 show rsvp interface detail on page 1274 show rsvp interface extensive on page 1274 show rsvp interface link-management on page 1275
Output Fields	<p>Table 49 on page 1272 lists the output fields for the show rsvp interface command. Output fields are listed in the approximate order in which they appear.</p>

Table 49: show rsvp interface Output Fields

Field Name	Field Description	Level of Output
RSVP interface	Number of interfaces on which RSVP is active. Each interface has one line of output.	All levels
Interface	Name of the interface.	All levels
Index	Index of the interface.	detail
State	State of the interface. <ul style="list-style-type: none"> • Disabled—No traffic engineering information is displayed. • Down—Interface is not operational. • Enabled—Displays traffic engineering information. • Up—Interface is operational. 	All levels
NoAuthentication	Interface does not support RSVP authentication.	detail
NoAggregate	Interface does not support refresh reduction.	detail
NoReliable	Interface does not support refresh reduction message ID extension.	detail
NoLinkProtection	Interface does not support link protection.	detail
HelloInterval	Frequency at which RSVP hellos are sent on this interface (in seconds).	detail
Address	IP address of the local interface.	detail
Active control channel	Next-hop link address to transmit messages.	None specified
TELink	Traffic-engineered links that are managed by the peer they are associated with.	None specified
Active resv	Number of reservations that are actively reserving bandwidth on the interface.	All levels
PreemptionCnt	Number of times an RSVP session was preempted on this interface.	detail
Update threshold	Percentage change in reserved bandwidth to trigger an IGP update.	detail
Subscription	User-configured subscription factor.	All levels
bc number	Bandwidth allocated for the specified bandwidth constraint.	extensive
ct number	Bandwidth allocated for the specified class type.	extensive
Static BW	Total interface bandwidth, in bps.	All levels
Available BW	Amount of bandwidth that RSVP is allowed to reserve, in bps. It is equal to (static bandwidth * subscription factor).	al levels

Table 49: show rsvp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Reserved BW	Currently reserved bandwidth, in bps.	All levels
SoftPreemptionCnt	Number of times a soft preemption occurred on this interface. This number is not included in the PreemptionCnt value.	detail
Overbooked BW	Currently overbooked bandwidth, in bps, by class type (ct0 through ct3).	detail
Highwater mark	Highest bandwidth that has ever been reserved on this interface, in bps.	brief
PacketType	Type of RSVP packet.	detail
Total Sent	Total number of packets sent.	detail
Total Received	Total number of packets received since RSVP was enabled.	detail
Last 5 seconds Sent	Number of packets sent in the last 5 seconds.	detail
Last 5 seconds Received	Number of packets received in the last 5 seconds.	detail
Path	Statistics about Path messages, which are sent from the RSVP sender along the data paths and store path state information in each node along the path.	detail
PathErr	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.	detail
PathTear	Statistics about PathTear messages, which remove path states and dependent reservation states in any routers along a path.	detail
Resv	Statistics about Resv messages, which are sent from the RSVP receiver along the data paths and store reservation state information in each node along the path.	detail
ResvErr	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.	detail
ResvTear	Statistics about ResvTear messages, which remove reservation states along a path.	detail
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.	detail
Ack	Acknowledge message for refresh reductions.	detail
Srefresh	Summary refresh messages.	detail
EndtoEnd RSVP	Statistics for the number of end-to-end RSVP messages sent.	detail
Queue	CoS transmit queue number and its associated forwarding class designation.	extensive

Table 49: show rsvp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
TxRate	Configured bandwidth in Mbps and configured bandwidth as a percentage of the specified queue.	extensive
Priority	Weight of the queue relative to other configured queues, in percentage.	extensive
queue-priority-value	Low, High, None, or Exact. None indicates no rate limiting. Exact indicates the queue transmits at the configured rate only.	extensive

Sample Output

show rsvp interface brief

```

user@host> show rsvp interface brief
RSVP interface: 1 active

```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
de0.0	Up	1	23%	10Mbps	989.992kbps	1.31Mbps	1.31Mbps

show rsvp interface detail

```

user@host> show rsvp interface detail
so-0/1/1.0 Index 6, State: Ena/Up
NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
HelloInterval 3(second)
Address 192.168.207.29, 10.255.245.194
ActiveResv 0, PreemptionCnt 0, Update threshold 10%
Subscription 100%, StaticBW 155.52Mbps, AvailableBW 155.52Mbps
ReservedBW [0] 155Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
SoftPreemptionCnt1
OverbookedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 155Mbps[5] 0bps[6] 0bps[7] 0bps
PacketType
Total
Last 5 seconds

```

	Sent	Received	Sent	Received
Path	16	0	1	0
PathErr	0	0	0	0
PathTear	1	0	0	0
Resv	0	11	0	1
ResvErr	0	0	0	0
ResvTear	0	0	0	0
Hello	66	67	1	1
Ack	0	0	0	0
Srefresh	0	0	0	0
EndtoEnd RSVP	0	0	0	0

...

show rsvp interface extensive

```

user@host> show rsvp interface extensive
so-1/0/0.0 Index 72, State Ena/Up
NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
HelloInterval 9(second)
Address 192.168.213.22, 10.255.240.175
ActiveResv 1, PreemptionCnt 0, Update threshold 10%
Subscription 100%,
bc0 = (ct0+ct1+ct2+ct3), StaticBW 622.08Mbps
bc1 = (ct1+ct2+ct3), StaticBW 466.56Mbps

```



```

bc2 = (ct2+ct3), StaticBW 311.04Mbps
bc3 = ct3, StaticBW 155.52Mbps
ct0: StaticBW 155.52Mbps, AvailableBW 522.08Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
ct1: StaticBW 155.52Mbps, AvailableBW 366.56Mbps
ReservedBW [0] 100Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps

ct2: StaticBW 155.52Mbps, AvailableBW 311.04Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
ct3: StaticBW 155.52Mbps, AvailableBW 155.52Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
Queue          TxRate          Priority Exact
0              155.52Mbps          25%     Low
1              155.52Mbps          25%     Low
2              155.52Mbps          25%     Low
3              155.52Mbps          25%     Low

```

show rsvp interface link-management

```

user@host> show rsvp interface link-management
RSVP interface: 2 active
PEER-C State: Up
Active Control Channel: so-0/1/0.0

TElink: TElnk1, Link ID: 37811
ActiveResv 0, PreemptionCnt 0
StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

TElink: TElnk2, Link ID: 37808
ActiveResv 1, PreemptionCnt 0
StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

PEER-B State: Up
Active Control Channel: so-1/0/0.0

TElink: TElnkAB1, Link ID: 1598
ActiveResv 0, PreemptionCnt 0
StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps

TElink: TElnkAB2, Link ID: 1597
ActiveResv 0, PreemptionCnt 0
StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps

```

show rsvp neighbor

List of Syntax	Syntax on page 1276 Syntax (EX Series Switches) on page 1276
Syntax	<pre>show rsvp neighbor <brief detail> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	<pre>show rsvp neighbor <brief detail></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>instance <i>instance-name</i> option added in Junos OS Release 15.1.</p>
Description	Display Resource Reservation Protocol (RSVP) neighbors that were discovered dynamically during the exchange of RSVP packets.
Options	<p>none—Display standard information about RSVP neighbors.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display the RSVP neighbor information for the specified instance. If <i>instance-name</i> is omitted, RSVP neighbor information is displayed for the master instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show rsvp neighbor on page 1280 show rsvp neighbor detail on page 1280
Output Fields	Table 50 on page 1276 lists the output fields for the show rsvp neighbor command. Output fields are listed in the approximate order in which they appear.

Table 50: show rsvp neighbor Output Fields

Field Name	Field Description	Level of Output
RSVP neighbor	Number of neighbors that the routing device has learned of. Each neighbor has one line of output.	All levels
via	Name of the interface where the neighbor has been detected. In the case of generalized MPLS (GMPLS) LSPs, the name of the peer where the neighbor has been detected.	detail
Address	Address of a learned neighbor.	All levels

Table 50: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Idle	Length of time the neighbor has been idle, in seconds.	All levels
Up/Dn	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.	All levels
Up cnt and Down cnt	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.	detail
status	State of the RSVP neighbor: <ul style="list-style-type: none"> • Up—Routing device can detect RSVP Hello messages from the neighbor. • Down—Routing device has received one of the following indications: <ul style="list-style-type: none"> • Communication failure from the neighbor. • Communication from IGP that the neighbor is unavailable. • Change in the sequence numbers in the RSVP Hello messages sent by the neighbor. • Restarting—RSVP neighbor is unavailable and might be restarting. The neighbor remains in this state until it has restarted or is declared dead. This state is possible only when graceful restart is enabled. • Restarted—RSVP neighbor has restarted and is undergoing state recovery (graceful restart) procedures. • Dead—Routing device has lost all communication with the RSVP neighbor. Any RSVP sessions with that neighbor are torn down. 	detail
LastChange	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is hh:mm:ss .	All levels
Last changed time	Time elapsed since the neighbor state changed either from up to down or from down to up.	detail
HelloInt	Frequency at which RSVP hellos are sent on this interface (in seconds).	All levels
HelloTx/Rx	Number of hello packets sent to and received from the neighbor.	All levels
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.	detail
Message received	Number of Path and Resv messages that this routing device has received from the neighbor.	detail
Remote Instance	Identification provided by the remote routing device during Hello message exchange.	detail

Table 50: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local Instance	Identification sent to the remote routing device during Hello message exchange.	detail
Refresh reduction	<p>Measure of processing overhead requests of refresh messages. Refresh reduction extensions improve routing device performance by reducing the process overhead, thus increasing the number of LSPs a routing device can support. Refresh reduction can have the following values:</p> <ul style="list-style-type: none"> • operational—All four RSVP refresh reduction extensions—message ack, bundling, summary refresh, and staged refresh timer—are functional between the two neighboring routing devices. For a detailed explanation of these extensions, see RFC 2961. • incomplete—Some RSVP refresh reduction extensions are functional between the two neighboring routing devices. • no operational—Either the refresh reduction feature has been turned off, or the remote routing device cannot support the refresh reduction extensions. 	detail
Remote end	<p>Neighboring routing device's status with regard to refresh reduction:</p> <ul style="list-style-type: none"> • enabled—Remote routing device has requested refresh reduction during RSVP message exchanges. • disabled—Remote routing device does not require refresh reduction. 	detail
Ack-extension	<p>An RSVP refresh reduction extension:</p> <ul style="list-style-type: none"> • enabled—Both local and remote routing devices support the ack-extension (RFC 2961). • disabled—Remote routing device does not support the ack-extension. 	detail
Link protection	<p>Status of the MPLS fast reroute mechanism that protects traffic from link failure:</p> <ul style="list-style-type: none"> • enabled—Link protection feature has been turned on, protecting the neighbor with a bypass LSP. • disabled—No link protection feature has been enabled for this neighbor. 	detail
LSP name	Name of the bypass LSP.	detail
Bypass LSP	<p>Status of the bypass LSP. It can have the following values:</p> <ul style="list-style-type: none"> • does not exist—Bypass LSP is not available. • connecting—Routing device is in the process of establishing a bypass LSP, and the LSP is not available for link protection at the moment. • operational—Bypass LSP is up and running. • down—Bypass LSP has gone down, with the most probable cause a node or a link failure on the bypass path. 	detail
Backup routes	Number of user LSPs (or routes) that are being protected by a bypass LSP (before link failure).	detail
Backup LSPs	Number of LSPs that have been temporarily established to maintain traffic by refreshing the downstream LSPs during link failure (not a one-to-one correspondence).	detail

Table 50: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Bypass explicit route	Explicit route object's (ERO) path that is taken by the bypass LSP.	detail
Restart time	Length of time a neighbor waits to receive a Hello from the restarting node before declaring the node dead and deleting the states (in milliseconds).	detail
Recovery time	Length of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds). Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.	detail

Sample Output

show rsvp neighbor

```
user@host> show rsvp neighbor
RSVP neighbor: 2 learned
Address          Idle Up/Dn LastChange HelloInt HelloTx/Rx
192.168.207.203   0 3/2    13:01      3   366/349
192.168.207.207   0 1/0    22:49      3   448/448
```

show rsvp neighbor detail

```
user@host> show rsvp neighbor detail
RSVP neighbor: 2 learned
Address: 192.168.207.203   via: ecstasy1 status: Up
  Last changed time: 28:47, Idle: 0 sec, Up cnt: 3, Down cnt: 2
  Message received: 632
  Hello: sent 673, received 656, interval 3 sec
  Remote instance: 0x6432838a, Local instance: 0x74b72e36
  Refresh reduction: operational
    Remote end: enabled, Ack-extension: enabled
  Link protection: enabled
    LSP name: Bypass_to_192.168.207.203
    Bypass LSP: operational, Backup routes: 1, Backup LSPs: 0
    Bypass explicit route: 192.168.207.207 192.168.207.224
  Restart time: 60000 msec, Recovery time: 0 msec
```

show rsvp session

List of Syntax	Syntax on page 1281 Syntax (EX and QFX Series Switches) on page 1281
Syntax	<pre>show rsvp session <brief detail extensive terse> <bidirectional unidirectional> <bypass> <down up> <externally-provisioned> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)> <lsp-type> <name <i>session-name</i>> <p2mp> <session-type> <statistics> <te-link <i>te-link</i>></pre>
Syntax (EX and QFX Series Switches)	<pre>show rsvp session <brief detail extensive terse> <bidirectional unidirectional> <bypass> <down up> <externally-provisioned> <interface <i>interface-name</i>> <lsp-type> <name <i>session-name</i>> <p2mp> <session-type> <statistics> <te-link <i>te-link</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>externally-provisioned option added in Junos OS Release 13.3.</p> <p>Command introduced in Junos OS Release 13.2X51-D15 for QFX Series.</p> <p>instance <i>instance-name</i> option added in Junos OS Release 15.1.</p>
Description	Display information about Resource Reservation Protocol (RSVP) sessions.
Options	<p>none—Display standard information about all RSVP sessions.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>bidirectional unidirectional—(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.</p> <p>bypass—(Optional) Display RSVP sessions for bypass LSPs.</p> <p>down up—(Optional) Display only LSPs that are inactive or active, respectively.</p>

externally-provisioned—(Optional) Display the LSPs that are generated dynamically and provisioned by an external Path Computation Element (PCE).

instance *instance-name*—(Optional) Display RSVP sessions for the specified instance. If *instance-name* is omitted, RSVP session information is displayed for the master instance.

interface *interface-name*—(Optional) Display RSVP sessions for the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-type—(Optional) Display information about RSVP sessions with regard to LSPs:

- **bypass**—Sessions used for bypass LSPs.
- **lsp**—Sessions used to set up LSPs.
- **nolsp**—Sessions not used to set up LSPs.

name *session-name*—(Optional) Display information about the named session.

p2mp—(Optional) Display point-to-multipoint information.

session-type—(Optional) Display information about a particular session type:

- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that transit through this routing device.

statistics—(Optional) Display packet statistics.

te-link *te-link*—(Optional) Display sessions with reservations on the specified TE link.

**Required Privilege
Level**

view

**Related
Documentation**

- [clear rsvp session on page 1260](#)

List of Sample Output

[show rsvp session on page 1286](#)
[show rsvp session statistics on page 1286](#)
[show rsvp session detail on page 1287](#)
[show rsvp session detail \(When Egress Protection is in Standby Mode\) on page 1287](#)
[show rsvp session detail \(When Egress Protection is in Effect During a Local Repair\) on page 1287](#)
[show rsvp session detail \(Path MTU Output Field\) on page 1288](#)
[show rsvp session detail \(GMPLS\) on page 1288](#)
[show rsvp session extensive on page 1288](#)
[show rsvp session p2mp \(Ingress Router\) on page 1289](#)
[show rsvp session p2mp \(Transit Router\) on page 1289](#)

Output Fields Table 51 on page 1283 describes the output fields for the **show rsvp session** command. Output fields are listed in the approximate order in which they appear.

Table 51: show rsvp session Output Fields

Field Name	Field Description	Level of Output
Ingress RSVP	Information about ingress RSVP sessions.	detail
Ingress RSVP	Information about ingress RSVP sessions. Each session has one line of output.	All levels
Egress RSVP	Information about egress RSVP sessions.	All levels
Transit RSVP	Information about the transit RSVP sessions.	All levels
P2MP name	(Appears only when the p2mp option is specified). Name of the point-to-multipoint LSP path.	All levels
P2MP branch count	(Appears only when the p2mp option is specified). Number of LSPs receiving packets from the point-to-multipoint LSP.	All levels
To	Destination (egress routing device) of the session.	All levels
From	Source (ingress routing device) of the session.	All levels
State	State of the path: Up , Down , or AdminDn . AdminDn indicates that the LSP is being taken down gracefully.	All levels
Address	Destination (egress routing device) of the LSP.	detail
From	Source (ingress routing device) of the session.	detail
LSPstate	State of the LSP that is being handled by this RSVP session. It can be either Up , Dn (down), or AdminDn . AdminDn indicates that the LSP is being taken down gracefully.	brief detail
Rt	Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	brief
Active Route	Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table (inet.0). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (mpls.0).	detail
LSPname	Name of the LSP.	brief detail
LSPpath	Indicates whether the RSVP session is for the primary or secondary LSP path. LSPpath can be either primary or secondary and can be displayed on the ingress, egress, and transit routing devices. LSPpath can also indicate when a graceful LSP deletion has been triggered.	detail

Table 51: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Bypass	(Egress routing device) Destination address for the bypass LSP.	detail
Bidir	(When LSP is bidirectional) LSP will allow data to travel in both directions between GMPLS devices.	detail
Bidirectional	(When LSP is bidirectional) LSP will allow data to travel both ways between GMPLS devices.	detail
Upstream label in	(When LSP is bidirectional) Incoming label for reverse direction traffic for this LSP.	detail
Upstream label out	(When LSP is bidirectional) Outgoing label for reverse direction traffic for this LSP.	detail
Recovery label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	detail
Recovery label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv messages that is returned.	detail
Suggested label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	detail
Suggested label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv message that is returned.	detail
Resv style or Style	RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).	brief detail
Label in	Incoming label for this LSP.	brief detail
Label out	Outgoing label for this LSP.	brief detail
Time left	Number of seconds remaining in the lifetime of the reservation.	brief detail
Since	Date and time when the RSVP session was initiated.	detail
Tspec	Sender's traffic specification, which describes the sender's traffic parameters.	detail
DiffServ info	Indicates whether the LSP is a multiclass LSP (multiclass diffServ-TE LSP) or a Differentiated-Services-aware traffic engineering LSP (diffServ-TE LSP).	detail
bandwidth	Bandwidth for each class type (ct0 , ct1 , ct2 , or ct3).	detail
Port number	Protocol ID and sender/receiver port used in this RSVP session.	detail
Attrib flags	Non-PHP indicates that ultimate hop popping has been requested by the LSP using this RSVP session	extensive

Table 51: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
FastReroute desired	Fast reroute has been requested by the ingress routing device.	detail
Soft preemption desired	Soft preemption has been requested by the ingress routing device.	detail
FastReroute desired	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Fast reroute (one-to-one backup) has been requested by the ingress routing device.	detail extensive
Link protection desired	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Link protection (many-to-one backup) has been requested by the ingress routing device.	detail extensive
Node/Link protection desired	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Node and link protection (many-to-one backup) has been requested by the ingress routing device.	detail extensive
Type	<p>LSP type:</p> <ul style="list-style-type: none"> • Link protected LSP—LSP has been protected by link protection at the outgoing interface. The name of the bypass used is also listed here (extensive). • Node/Link protected LSP—LSP has been protected by node and link protection at the outgoing interface. The name of the bypass used is also listed here (extensive). • Protection down—LSP is not currently protected. • Bypass LSP—LSP that is used to protect one or more user LSPs in case of link failure. • Backup LSP at Point-of-Local-Repair (PLR)—LSP that has been temporarily established to protect a user LSP at the ingress of a failed link. • Backup LSP at Merge Point (MP)—LSP that has been temporarily established to protect a user LSP at the egress of a failed link. 	detail extensive
New bypass	New bypass (the bypass name is also displayed) has been activated to protect the LSP.	extensive
Link protection up, using <i>bypass-name</i>	Link protection (the bypass name is also displayed) has been activated for the LSP.	extensive
Creating backup LSP, link down	A link down event occurred, and traffic is being switched over to the bypass LSP.	extensive
Deleting backup LSP, protected LSP restored	Link has come back up and the LSP has been restored. Because the backup LSP is no longer needed, it is deleted.	extensive
Path mtu	Displays the value of the path MTU received from the network (through signaling) and the value used for forwarding. This value is only displayed on ingress routing devices with the allow-fragmentation statement configured at the [edit protocols mpls path-mtu] hierarchy level. If there is a detour LSP, the path MTU for the detour is also displayed.	detail

Table 51: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Egress protection PLR as protector	RSVP state on the Protector or the point-of-local-repair (PLR) routing device: <ul style="list-style-type: none"> Active— Egress protection is available at the Protector or the PLR routing device. In Use— Local repair has been completed. 	detail
PATH rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor.	detail
Adspec	MTU signaled from the ingress routing device to the egress routing device by means of the adspec object.	detail
PATH sentto	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor (or peer-name in the GMPLS LSP case), and number of packets sent to the downstream routing device.	detail
Explct route	Explicit route for the session. Normally this value will be the same as that of record route. Differences indicate that path rerouting has occurred, typically during fast reroute.	detail
Record route	Recorded route for the session, taken from the record route object. Normally this value will be the same as that of explct route. Differences indicate that path rerouting has occurred, typically during fast reroute.	detail

Sample Output

show rsvp session

```

user@host> show rsvp session
Ingress RSVP: 1 sessions
To          From          State  Rt  Style Labelin Labelout LSPName
10.255.245.214 10.255.245.212 AdminDn 0 1 FF - 22293 LSP Bidir
Total 1 displayed, Up 1, Down 0

Egress RSVP: 2 sessions
To          From          State  Rt  Style Labelin Labelout LSPName
10.255.245.194 10.255.245.195 Up    0 1 FF 39811 - Gpro3-ba Bidir
10.255.245.194 10.255.245.195 Up    0 1 FF 3 - pro3-ba
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To          From          State  Rt  Style Labelin Labelout LSPName
10.255.245.198 10.255.245.197 Up    0 1 SE 100000 3 pro3-de
Total 1 displayed, Up 1, Down 0

```

show rsvp session statistics

```

user@host> show rsvp session statistics
Ingress RSVP: 2 sessions
To          From          State  Packets  Bytes  LSPName
10.255.245.24 10.255.245.22 Up    0        0  pro3-bd
10.255.245.24 10.255.245.22 Up   44868 2333136 pro3-bd-2

```

```

Total 2 displayed, Up 2, Down 0
Egress RSVP: 2 sessions
To           From           State   Packets   Bytes   LSPname
10.255.245.22 10.255.245.24   Up      0         0      pro3-db
10.255.245.22 10.255.245.24   Up      0         0      pro3-db-2
Total 2 displayed, Up 2, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

show rsvp session detail

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPname: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left: -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt

```

show rsvp session detail (When Egress Protection is in Standby Mode)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPname: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left: -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  Egress protection PLR as protector: Active
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt

```

show rsvp session detail (When Egress Protection is in Effect During a Local Repair)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Down, ActiveRoute: 0
  LSPname: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left: -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  Egress protection PLR as protector: In Use
  PATH rcvfrom: localclient

```

```

Adspec: sent MTU 1500
PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt

```

show rsvp session detail (Path MTU Output Field)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
10.255.245.3
  From: 10.255.245.5, LSPstate: Up, ActiveRoute: 3
  LSPname: to-c, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100432
  Resv style: 1 FF, Label in: -, Label out: 100432
  Time left: -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
  Port number: sender 1 receiver 57843 protocol 0
  FastReroute desired
  PATH rcvfrom: localclient
  Adspec: sent MTU 4470
  Path mtu: received 4470, using 4458 for forwarding
  PATH sentto: 192.168.37.89 (so-0/2/3.0) 11 pkts
  RESV rcvfrom: 192.168.37.89 (so-0/2/3.0) 10 pkts
  Explct route: 192.168.37.89
  Record route: <self> 192.168.37.89 192.168.37.87
    Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
    Detour adspec: sent MTU 1512
    Path mtu: received 1512, using 1500 for forwarding

```

show rsvp session detail (GMPLS)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
192.168.4.1
  From: 192.168.1.1, LSPstate: Dn, ActiveRoute: 0
  LSPname: gmpls-r1-to-r3, LSPpath: Primary
  Bidirectional, Upstream label in: 21253, Upstream label out: -
  Suggested label received: -, Suggested label sent: 21253
  Recovery label received: -, Recovery label sent: -
  Resv style: 0 -, Label in: -, Label out: -
  Time left: -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
  Port number: sender 2 receiver 46115 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH MTU: received 0
  PATH sentto: 10.35.1.5 (so-0/2/3.0) 11 pkts
  Explct route: 100.100.100.100 93.93.93.93
  Record route: <self> 100.100.100.100 93.93.93.93
  Total 1 displayed, Up 0, Down 1
  Egress RSVP: 0 sessions
  Total 0 displayed, Up 0, Down 0
  Transit RSVP: 0 sessions
  Total 0 displayed, Up 0, Down 0

```

show rsvp session extensive

```

user@host> show rsvp session extensive
Ingress RSVP: 1 sessions

192.168.0.4
  From: 192.168.0.5, LSPstate: Up, ActiveRoute: 0

```

```

LSPname: E-D, LSPpath: Primary
LSPtype: Static Configured
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 299808
Resv style: 1 FF, Label in: -, Label out: 299808
Time left: -, Since: Thu Sep 20 15:54:20 2012
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 2 receiver 61576 protocol 0
Attrib flags: Non-PHP
PATH rcvfrom: localclient
Adspec: sent MTU 1500
Path MTU: received 1500
PATH sentto: 10.0.0.18 (lt-1/2/0.17) 41 pkts
RESV rcvfrom: 10.0.0.18 (lt-1/2/0.17) 40 pkts
Explct route: 10.0.0.18 10.0.0.22
Record route: <self> 10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

```

Egress RSVP: 1 sessions

192.168.0.5

```

From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
LSPname: E-D, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 140, Since: Thu Sep 20 15:52:10 2012
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 49601 protocol 0
PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 44 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

```

Transit RSVP: 0 sessions

Total 0 displayed, Up 0, Down 0

show rsvp session p2mp (Ingress Router)

```

user@host> show rsvp session p2mp
Ingress RSVP: 3 sessions
P2MP name: test, P2MP branch count: 1
To      From      State  Rt Style Labelin Labelout LSPname
10.255.10.95 10.255.10.2  Up    0  1 SE  -        3 to-pe1
P2MP name: test2, P2MP branch count: 2
To      From      State  Rt Style Labelin Labelout LSPname
10.255.10.23 10.255.10.2  Up    0  1 SE  -        299776 to-pe3
10.255.10.16 10.255.10.2  Up    0  1 SE  -        299776 to-pe4
Total 3 displayed, Up 3, Down 0

```

Egress RSVP: 0 sessions

Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions

Total 0 displayed, Up 0, Down 0

show rsvp session p2mp (Transit Router)

```

user@host> show rsvp session p2mp

```

Ingress RSVP: 1 sessions

P2MP name: test, P2MP branch count: 1

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.255.10.23	10.255.10.95	Up	0	1 SE	-	299792	to-pe2

Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions

P2MP name: test, P2MP branch count: 1

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.255.10.95	10.255.10.2	Up	0	1 SE	3	-	to-pe1

Total 1 displayed, Up 1, Down 0

Transit RSVP: 2 sessions

P2MP name: test2, P2MP branch count: 2

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.255.10.23	10.255.10.2	Up	0	1 SE	299776	299808	to-pe3
10.255.10.16	10.255.10.2	Up	0	1 SE	299776	299856	to-pe4

Total 2 displayed, Up 2, Down 0

show rsvp statistics

List of Syntax	Syntax on page 1291 Syntax (EX Series Switches) on page 1291
Syntax	<pre>show rsvp statistics <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	show rsvp statistics
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>instance <i>instance-name</i> option added in Junos OS Release 15.1.</p>
Description	Display Resource Reservation Protocol (RSVP) packet and error statistics.
Options	<p>none—Display RSVP packet and error statistics.</p> <p>instance <i>instance-name</i>—(Optional) Display RSVP packet and error statistics for the specified instance. If <i>instance-name</i> is omitted, RSVP statistics is displayed for the master instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear rsvp statistics on page 1262
List of Sample Output	show rsvp statistics on page 1294
Output Fields	<p>Table 52 on page 1291 describes the output fields for the show rsvp statistics command. Output fields are listed in the approximate order in which they appear.</p>

Table 52: show rsvp statistics Output Fields

Field Name	Field Description
Packet Type	Statistics about different RSVP messages.
Total Sent	Total number of packets sent since RSVP was enabled.
Total Received	Total number of packets received since RSVP was enabled.
Last 5 seconds Sent	Total number of packets sent in the last 5 seconds.
Last 5 seconds Received	Number of packets received in the last 5 seconds.

Table 52: show rsvp statistics Output Fields (*continued*)

Field Name	Field Description
Path	Statistics about Path messages, which are sent from the RSVP sender along the data paths and which store path state information in each node along the path.
PathErr	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.
PathTear	Statistics about PathTear messages, which remove path states and dependent reservation states in any routing devices along a path.
Resv FF	Statistics about fixed-filter reservation style messages, which consist of distinct reservations among explicit senders.
Resv WF	Statistics about wildcard-filter reservation style messages, which consist of shared reservations among wildcard senders.
Res SE	Statistics about shared-explicit reservation style messages, which consist of shared reservations among explicit senders.
ResvErr	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.
ResvTear	Statistics about ResvTear messages, which remove reservation states along a path.
ResvConf	Statistics about ResvConfirm messages, which are responses to confirm a reservation request.
Ack	Acknowledge message for refresh reductions.
SRefresh	Summary refresh messages.
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.
EndtoEnd RSVP	Statistics for the number of End-to-end RSVP messages.
Errors	Statistics about errored RSVP packets.
Rcv pkt bad length	The packet was not processed because its length is inappropriate.
Rcv pkt unknown type	The packet is not one of the well-known RSVP types, as defined in RFC 2205, <i>Resource ReSerVation Protocol (RSVP)</i> .
Rcv pkt bad version	The packet is not an RSVP version 1 packet.
Rcv pkt auth fail	The packet failed authentication checks.
Rcv pkt bad checksum	The RSVP checksum check failed.
Rcv pkt bad format	General packet processing failed because the packet was badly formed.
Memory allocation fail	An internal resource failure occurred.

Table 52: show rsvp statistics Output Fields (*continued*)

Field Name	Field Description
No path information	A reservation was received, but no sender is active.
Resv style conflict	The same session contains inconsistent reservation styles.
Port conflict	There were inconsistent port numbers for the same session.
Resv no interface	An interface for the receive reservation packets cannot be located.
PathErr to client	Number of PathErr packets delivered to the local client.
ResvErr to client	Number of ResvErr packets delivered to the local client.
Path timeout	Number of times the sender timed out because the path was removed.
Resv timeout	Number of times the receiver timed out because the reservation was removed.
Message out-of-order	Records the number of RSVP incoming messages that are considered out of order. This is detected from the message ID object's sequence number.
Unknown ack msg	A neighboring routing device replies with an ACK object that contains an unknown message ID. This can indicate a message ID handshake problem. For example, a router receives an ACK for message IDs 1, 2, and 3. However, it only has state for message IDs 1 and 3. The router increments the unknown ack counter by 1.
Recv nack	If a neighboring router receives an unknown message ID in an RSVP refresh message, the router sends a Resv nack message back to the sender. This can happen if that neighbor has been rebooted. For this case, the router sends a regular RSVP refresh message to recover the state and start the message-ID handshake process again.
Recv duplicated msg-id	Number of times the same message ID is used by two different RSVP messages. This duplication is usually caused when a neighboring routing device restarts.
No TE-link to rcv Hop	Counter of packets discarded because a TE link was not found.
Rcv pkt disabled interface	Number of RSVP packets received on an interface that is not enabled for RSVP.
Transmit buffer full	Number of times the buffer for assembling an outgoing RSVP message was not large enough.
Transmit failure	Number of times the RSVP task failed to send out a packet.
Receive failure	Number of times the RSVP task failed to read an incoming packet.
P2MP RESV discarded by appl	Number of Resv messages discarded because the MPLS label is not valid for the P2MP LSP application.
Rate limit	Number of RSVP packets dropped due to rate limiting.

Table 52: show rsvp statistics Output Fields (*continued*)

Field Name	Field Description
Err msg loop detected	Number of RSVP error messages that have looped back to their originator. This is detected by checking the error node address in the ERROR_SPEC object.

Sample Output

show rsvp statistics

```

user@host> show rsvp statistics
  PacketType          Sent      Received      Last 5 seconds
                                     Sent      Received
Path                  355         408           0           0
PathErr                2          13           0           0
PathTear              101        139           0           0
Resv FF                0           0           0           0
Resv WF                0           0           0           0
Resv SE               419        225           0           0
ResvErr                0           0           0           0
ResvTear               0          13           0           0
ResvConf               0           0           0           0
Ack                   682       1414           0           0
SRefresh              395198     236030         5           2
Hello                 578809     578221         4           4
EndtoEnd RSVP         0           0           0           0

Errors                Total      Last 5 seconds
Rcv pkt bad length      0           0
Rcv pkt unknown type    0           0
Rcv pkt bad version     0           0
Rcv pkt auth fail       0           0
Rcv pkt bad checksum    0           0
Rcv pkt bad format      0           0
Memory allocation fail  0           0
No path information     10          0
Resv style conflict     0           0
Port conflict           0           0
Resv no interface       0           0
PathErr to client       38          0
ResvErr to client       0           0
Path timeout            8           0
Resv timeout            57          0
Message out-of-order    0           0
Unknown ack msg         2978        0
Recv nack               86          0
Recv duplicated msg-id   5           0
No TE-link to recv Hop  0           0
Rcv pkt disabled interface 0           0
Transmit buffer full    0           0
Transmit failure        0           0
Receive failure         0           0
P2MP RESV discarded by appl 0           0
Rate limit              306         0
Err msg loop detected    0           0

```

show rsvp version

List of Syntax	Syntax on page 1295 Syntax (EX Series Switches) on page 1295
Syntax	show rsvp version <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show rsvp version
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display information about the Resource Reservation Protocol (RSVP) protocol settings, such as the version of the RSVP software, the refresh timer and keep multiplier, and local RSVP graceful restart capabilities on a routing device.
Options	none —Display RSVP protocol settings. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show rsvp version on page 1296
Output Fields	Table 53 on page 1295 describes the output fields for the show rsvp version command. Output fields are listed in the approximate order in which they appear.

Table 53: show rsvp version Output Fields

Field Name	Field Description
Resource ReSerVation Protocol, version	RSVP software version.
RSVP protocol	Status of RSVP: Enabled or Disabled .
R(refresh timer)	Configured time interval used to generate periodic RSVP messages.
K(keep multiplier)	Number of RSVP messages that can be lost before an RSVP state is declared stale.
Preemption	Currently configured preemption capability: Aggressive , Disabled , or Normal . The default is Normal .
Soft-preemption cleanup	Time, in seconds, that an LSP is kept after it has been soft preempted. This is a global property of the RSVP protocol.
Graceful deleting timeout	Currently configured value for the graceful-deletion-timeout statement. The router that initiates the graceful deletion procedure for an RSVP session waits for the graceful deletion timeout interval to ensure that all routers along the path (especially the ingress and egress routers) have prepared for the LSP to be taken down.

Table 53: show rsvp version Output Fields (*continued*)

Field Name	Field Description
NSR Mode	Status of the nonstop active routing feature for RSVP on the restarting device: Disabled , Enabled/Master , or Enabled/Standby .
NSR State	<p>State of the nonstop active routing feature for RSVP on the restarting device.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> • Idle • TE-link sync complete • Neighbor sync complete • Path state sync complete • Resv state sync complete • Bypass sync complete • Init sync complete
Setup protection	Status of point-to-point and point-to-multipoint LSP setup protection configuration on the device: Enabled or Disabled
Graceful restart	Status of the graceful restart feature for RSVP on the restarting routing device: Enabled or Disabled .
Restart helper mode	Status of the helper mode feature: Enabled or Disabled . When this feature is enabled, the restarting routing device can help the neighbor with its RSVP restart procedures.
Maximum helper restart time	Number of milliseconds (ms) configured for the maximum helper restart time. The maximum helper restart time is the length of time the routing device waits before declaring that an RSVP neighbor attempting to restart gracefully is down.
Maximum helper recovery time	Number of milliseconds configured for the maximum helper recovery time. The maximum helper recovery time is the amount of time the routing device maintains the state of an RSVP neighbor attempting to restart gracefully.
Restart time	Number of milliseconds that a neighbor waits to receive a Hello message from the restarting node before declaring the node dead and deleting the states.
Recovery time	Number of milliseconds during which the restarting node attempts to recover its lost states with help from its neighbors. Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.
P2p transit LSP nexthop mode	Point-to-point transit LSP nexthop mode on PTX Series devices. The possible values are Chained or Unchained
P2mp transit LSP nexthop mode	Point-to-multipoint transit LSP nexthop mode on PTX Series devices. The possible values are Chained or Unchained

Sample Output

show rsvp version

```
user@host> show rsvp version
```

```
Resource ReSerVation Protocol, version 1. rfc2205
  RSVP protocol:           Enabled
  R(refresh timer):        30 seconds
  K(keep multiplier):      3
  Preemption:              Normal
  Soft-preemption cleanup:  30 seconds
  Graceful deletion timeout: 30 seconds
  NSR mode:                 Enabled/Master
  NSR state:                Init sync complete
  Setup protection:        Disabled
  Graceful restart:        Disabled
  Restart helper mode:      Enabled
  Maximum helper restart time: 20000 msec
  Maximum helper recovery time: 180000 msec
  Restart time:             0 msec
  P2p transit LSP nexthop mode: Unchained
  P2mp transit LSP nexthop mode: Unchained
```

traceroute mpls rsvp

Syntax	<code>traceroute mpls <rsvp> <i>lsp-name</i></code> <code><detail></code> <code><egress></code> <code><exp></code> <code><logical-system></code> <code><multipoint></code> <code><no-resolve></code> <code><retries></code> <code><source <i>source-address</i>></code> <code><ttl></code>
Release Information	Command introduced in Junos OS Release 9.2. <code>egress</code> , <code>multipoint</code> , and <code>ttl</code> options added in Junos OS Release 11.2.
Description	Trace route to a remote host for an MPLS LSP signaled by RSVP. Use traceroute mpls rsvp as a debugging tool to locate MPLS label-switched path (LSP) forwarding issues in a network. (Currently supported for IPv4 packets only.)
Options	<p><i>lsp-name</i>—Specify the name of the LSP to be traced.</p> <p>detail—(Optional) Display detailed output.</p> <p>egress—(Optional) Request that a specific point-to-multipoint egress node reply to the trace route. The trace route would follow the associated sub-LSP to the egress node.</p> <p>exp—(Optional) Specify the class of service to use when sending probes. The range of values is 0 through 7. The default value is 7.</p> <p>logical-system—(Optional) Specify the name of the logical system for the traceroute attempt.</p> <p>multipoint—(Optional) Perform a trace route on a point-to-multipoint LSP.</p> <p>no-resolve—(Optional) Specify not to resolve the hostname that corresponds to the IP address.</p> <p>retries—(Optional) Specify the number of times to resend probe. The range of values is 1 through 9. The default value is 3.</p> <p>source <i>source-address</i>—(Optional) Specify the source address of the outgoing traceroute packets.</p> <p>ttl—(Optional) Specify the number of hops to follow before forcing the trace route to quit.</p>
Required Privilege Level	network
List of Sample Output	traceroute mpls rsvp on page 1300 traceroute mpls rsvp detail on page 1300

[traceroute mpls rsvp multipoint \(branch node for sub-LSPs\) on page 1301](#)
[traceroute mpls rsvp multipoint \(single-hop sub-LSPs\) on page 1301](#)

Output Fields Table 54 on page 1299 describes the output fields for the **traceroute mpls rsvp *lsp-name*** and **traceroute mpls rsvp *lsp-name* detail** commands. Output fields are listed in the approximate order in which they appear.

Table 54: traceroute mpls rsvp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the traceroute mpls rsvp <i>lsp-name</i> command.	all levels
ttl	Time-to-live value of the labeled packet.	none specified
Label	MPLS label used to forward the packets along the LSP.	none specified
Protocol	Signaling protocol used. For this command, it is RSVP-TE.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is null.	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths). Displays Success if the trace to a hop is successful or Egress if the trace has reached the last router on the path.	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	detail
Parent	Address of the previous hop. Parent value for the first hop is null.	detail
Return Code	Return code for reporting the result of processing the echo request by the receiver.	detail
Sender timestamp	Displays the timestamp when the MPLS echo request is sent to the next hop.	detail
Receiver timestamp	Timestamp when the echo request from the previous hop is received and acknowledged with an echo response by the next hop.	detail
Response time	Time for the echo request to reach the receiver.	detail
MTU	Size of the largest packet that includes the label stack forwarded to the next hop.	detail

Table 54: traceroute mpls rsvp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is none.	detail
Label stack	Label stack used to forward the packet.	detail
Path	Displays the sub-lsp path number for this traceroute, the interface used, and the destination address.	all levels

Sample Output

traceroute mpls rsvp

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta
```

```
Probe options: retries 3, exp 7
```

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	299792	RSVP-TE	192.168.1.2	(null)	Success
2	299803	RSVP-TE	192.168.2.3	192.168.1.2	Success
3	3	RSVP-TE	192.168.3.4	192.168.2.3	Egress

```
Path 1 via ge-0/0/0.1 destination 127.0.0.64
```

traceroute mpls rsvp detail

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta detail
```

```
Probe options: retries 3, exp 7
```

```
Hop 192.168.1.2 Depth 1
```

```
Probe status: Success
```

```
Parent: (null)
```

```
Return code: Label-switched at stack-depth 1
```

```
Sender timestamp: 2008-04-17 09:35:27 EDT 400.88 msec
```

```
Receiver timestamp: 2008-04-17 09:35:27 EDT 427.87 msec
```

```
Response time: 26.99 msec
```

```
MTU: Unknown
```

```
Multipath type: IP bitmask
```

```
Address Range 1: 127.0.0.64 ~ 127.0.0.127
```

```
Label Stack:
```

```
Label 1 Value 299792 Protocol RSVP-TE
```

```
Hop 192.168.2.3 Depth 2
```

```
Probe status: Success
```

```
Parent: 192.168.1.2
```

```
Return code: Upstream interface index unknown label-switched at stack-depth
```

```
1
```

```
Sender timestamp: 2008-04-17 09:35:27 EDT 522.13 msec
```

```
Receiver timestamp: 2008-04-17 09:35:27 EDT 548.69 msec
```

```
Response time: 26.55 msec
```

```
MTU: 1518
```

```
Multipath type: IP bitmask
```

```
Address Range 1: 127.0.0.64 ~ 127.0.0.127
```

```
Label Stack:
```

```
Label 1 Value 299803 Protocol RSVP-TE
```

traceroute mpls rsvp multipoint (branch node for sub-LSPs)

The following traceroute output is for a point-to-multipoint LSP where the penultimate node is a branch node for the sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
Probe options: retries 3, exp 7
```

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	300000	RSVP-TE	81.1.2.2	(null)	Success
2	299968	RSVP-TE	81.2.3.3	81.1.2.2	Success
3	299952	RSVP-TE	81.3.4.4	81.2.3.3	Success
4	299920	RSVP-TE	81.4.6.6	81.3.4.4	Egress

Path 1 via lt-1/2/0.102 destination 127.0.0.64

ttl	Label	Protocol	Address	Previous Hop	Probe Status
4	299920	RSVP-TE	81.4.5.5	81.3.4.4	Egress

Path 2 via lt-1/2/0.102 destination 127.0.0.64

traceroute mpls rsvp multipoint (single-hop sub-LSPs)

The following traceroute output is for a point-to-multipoint LSP with multiple single-hop sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
Probe options: retries 3, exp 7
```

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	0	RSVP-TE	81.1.2.2	(null)	Egress

Path 1 via lt-1/2/0.102 destination 127.0.0.64

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	0	RSVP-TE	81.1.8.8	(null)	Egress

Path 2 via lt-1/2/0.108 destination 127.0.0.64

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	0	RSVP-TE	81.1.9.9	(null)	Egress

Path 3 via lt-1/2/0.109 destination 127.0.0.64

CHAPTER 31

LDP Operational Commands

- `clear ldp neighbor`
- `clear ldp session`
- `clear ldp statistics`
- `ping mpls ldp`
- `show ldp database`
- `show ldp fec-filters`
- `show ldp interface`
- `show ldp neighbor`
- `show ldp overview`
- `show ldp p2mp tunnel`
- `show ldp path`
- `show ldp route`
- `show ldp session`
- `show ldp statistics`
- `show ldp traffic-statistics`
- `show security keychain`
- `traceroute mpls ldp`

clear ldp neighbor

Syntax	clear ldp neighbor <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> < <i>neighbor</i> >
Description	Tear down Label Distribution Protocol (LDP) neighbor connections.
Options	none —Tear down connections with all LDP neighbors for all routing instances. instance <i>instance-name</i> —(Optional) Clear the LDP session for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>neighbor</i> —(Optional) Clear an LDP session for the specified neighbor (IP address) only.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show ldp neighbor on page 1322
List of Sample Output	clear ldp neighbor on page 1304
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ldp neighbor

```
user@host> clear ldp neighbor
```

clear ldp session

Syntax	clear ldp session <destination> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Clear Label Distribution Protocol (LDP) sessions.
Options	<p>none—Clear LDP sessions for all destinations for all routing instances.</p> <p>destination—(Optional) Clear an LDP session for the specified destination (IP address).</p> <p>instance <i>instance-name</i>—(Optional) Clear the LDP session for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show ldp session on page 1335
List of Sample Output	clear ldp session on page 1305
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ldp session

```
user@host> clear ldp session
```

clear ldp statistics

Syntax	<code>clear ldp statistics</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Release Information	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Set all Label Distribution Protocol (LDP) statistics to zero.
Options	none —Set all LDP statistics to zero for all routing instances. instance <i>instance-name</i> —(Optional) Clear the LDP session for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show ldp statistics on page 1341• show ldp traffic-statistics on page 1345
List of Sample Output	clear ldp statistics on page 1306
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ldp statistics

```
user@host> clear ldp statistics
```


ping mpls ldp

Syntax	<pre>ping mpls ldp fec <count count> <destination address> <detail> <exp forwarding-class> <instance routing-instance-name> <logical-system (all logical-system-name)> <p2mp root-addr ip-address lsp-id identifier> <size bytes> <source source-address> <sweep></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>size and sweep options introduced in Junos OS Release 9.6.</p> <p>instance option introduced in Junos OS Release 10.0.</p> <p>p2mp, root-address, and lsp-id options introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	<p>Check the operability of MPLS LDP-signaled label-switched path (LSP) connections. Type Ctrl+c to interrupt a ping mpls command.</p>
Options	<p>count count—(Optional) Number of ping requests to send. If count is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.</p> <p>destination address—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.</p> <p>detail—(Optional) Display detailed information about the echo requests sent and received.</p> <p>exp forwarding-class—(Optional) Value of the forwarding class for the MPLS ping packets.</p> <p>fec—Ping an LDP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.</p> <p>instance routing-instance-name—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on the specified logical system.</p> <p>p2mp root-addr ip-address lsp-id identifier—(Optional) Ping the end points of a point-to-multipoint LSP. Enter the IP address of the point-to-multipoint LSP root and the ID number of the point-to-multipoint LSP.</p> <p>size bytes—(Optional) Size of the LSP ping request packet (88 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller</p>

than the minimum size, an error message is displayed reminding you of the 88-byte minimum.

source source-address—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

You can configure the ping interval for the **ping mpls ldp** command by specifying a new time in seconds using the **lsp-ping-interval** statement at the **[edit protocols ldp oam]** hierarchy level. For more information, see the *MPLS Applications Feature Guide for Routing Devices*.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level network

List of Sample Output [ping mpls ldp fec count on page 1308](#)
[ping mpls ldp p2mp root-addr lsp-id on page 1308](#)

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls ldp fec count

```
user@host> ping mpls ldp 10.255.245.222 count 10
!!!xxx...x--- lsping statistics ---10 packets transmitted, 3 packets received,
70% packet loss 4 packets received with error status, not counted as received.
```

ping mpls ldp p2mp root-addr lsp-id

```
user@host> ping mpls ldp p2mp root-addr 10.1.1.1/32 lsp-id 1 count 1
Request for seq 1, to interface 71, no label stack.
Request for seq 1, to interface 70, label 299786
Reply for seq 1, egress 10.1.1.3, return code: Egress-ok, time: 18.936 ms
Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
```

```
Reply for seq 1, egress 10.1.1.4, return code: Egress-ok, time: 18.936 ms
  Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
  Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
Reply for seq 1, egress 10.1.1.5, return code: Egress-ok, time: 18.936 ms
  Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
  Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
```

show ldp database

Syntax	<code>show ldp database</code> <code><brief detail extensive></code> <code><inet l2circuit></code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><p2mp></code> <code><session <i>session</i>></code> <code><summary></code>
Release Information	Command introduced before Junos OS Release 7.4. summary option introduced in Junos OS Release 14.2.
Description	Display entries in the LDP database.
Options	<p>none—Display standard information about all entries in the LDP database for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>inet l2circuit—(Optional) Display only IPv4 or Layer 2 circuit bindings.</p> <p>instance <i>instance-name</i>—(Optional) Display routing instance information for the specified instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>p2mp—(Optional) Display point-to-multipoint binding information.</p> <p>session <i>session</i>—(Optional) Display database for the specified session only. <i>session</i> is the destination address of the LDP session.</p> <p>summary—(Optional)—Display summary output. This option displays the number of labels received and advertised for each LDP session.</p>
Required Privilege Level	view
List of Sample Output	show ldp database (master) on page 1313 show ldp database (standby) on page 1314 show ldp database l2circuit detail on page 1314 show ldp database l2circuit extensive on page 1315 show ldp database p2mp (master) on page 1315 show ldp database p2mp (standby) on page 1315 show ldp database p2mp (master) on page 1316 show ldp database p2mp (standby) on page 1316 show ldp database session on page 1316 show ldp database (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 1317

[show ldp database \(Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 1317](#)
[show ldp database summary on page 1318](#)

Output Fields [Table 55 on page 1311](#) describes the output fields for the **show ldp database** command. Output fields are listed in the approximate order in which they appear.

Table 55: show ldp database Output Fields

Field Name	Field Description	Level of Output
Input label database	Label received from the other router.	All levels
Output label database	Label advertised to the other router.	All levels
<i>session-identifier</i>	Session identifier, which includes the local and remote label space identifiers.	All levels
Labels received	Number of labels received from the other router.	All levels
Labels advertised	Number of labels advertised to the other router.	All levels.
Label	Label binding to a route prefix.	All levels

Table 55: show ldp database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Prefix	<p>Route prefix.</p> <p>It can be one of the following values:</p> <ul style="list-style-type: none"> • IP prefix. • Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured. • Layer 2 encapsulation type. <p>Layer 2 encapsulation types are displayed in the format L2CKT control word status encapsulation-type vc-number, for example, L2CKT CtlfWord FRAME RELAY VC 2</p> <ul style="list-style-type: none"> • control-word-status—Displays whether the use of the control word has been negotiated for this virtual circuit: <ul style="list-style-type: none"> • NoCtrlWord • CtrlWord • encapsulation-type—Encapsulation type: <ul style="list-style-type: none"> • FRAME RELAY • ATM AAL5 • ATM CELL • VLAN • ETHERNET • CISCO_HDLC • PPP • VC number—Virtual circuit number. It can have any numeric value. • (Stale)—When you display the LDP database for the neighbor of a restarting router, the bindings learned from the restarting neighbor are displayed as (Stale). Stale bindings are deleted if they are not refreshed within the recovery time. 	All levels
MTU	MTU of the Layer 2 circuit. MTU is displayed for all encapsulation types except ATM cell encapsulations.	detail
VCCV Control Channel types	<p>Virtual Circuit Connection Verification (VCCV) control channel types.</p> <ul style="list-style-type: none"> • MPLS router alert label • MPLS PW label with TTL=1 	extensive
VCCV Control Verification types	The only valid VCCV control verification type is LSP ping .	extensive
TDM payload size	Size of the Time Division Multiplex (TDM) payload.	All levels
TDM bitrate	Bit rate for the TDM traffic.	All levels
Requested VLAN ID	(VLANs) VLAN identifier of the Layer 2 circuit.	detail
Cell bundle size	(ATM cell encapsulations) Maximum number of cells that the Layer 2 circuit can receive in a packet.	detail

Table 55: show ldp database Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	State of the label binding: <ul style="list-style-type: none"> Active—Label binding has been installed and distributed appropriately. A label binding is almost always in this state. New—New label that has not yet been distributed. <ul style="list-style-type: none"> MapRcv—Waiting to receive a label mapping message. MapSend—Waiting to send a label mapping message. RelRcv—Waiting to receive a label release message. RelRsnd—Waiting to receive a label release message before resending label mapping message. RelSend—Waiting to send a label release message. ReqSend—Waiting to send a label request message. W/dSend—Waiting to send a label withdrawal message. 	detail
Age	Time elapsed since the binding was created.	detail

Sample Output

show ldp database (master)

```

user@host> show ldp database extensive
Input label database, 10.255.107.232:0--10.255.107.236:0
  Label Prefix
  299840 10.255.107.232/32
           State: Active
           Age: 9:35
           Entropy Label Capability: No
           3 10.255.107.236/32
           State: Active
           Age: 9:35
           Entropy Label Capability: No
  299776 L2CKT CtrlWord VLAN VC 100
           MTU: 1500 Requested VLAN ID: 600 Flow Label T Bit: 1 Flow Label R
           Bit: 1
           State: Active
           Age: 9:35
           Entropy Label Capability: No
           VCCV Control Channel types:
             PWE3 control word
             MPLS router alert label
             MPLS PW label with TTL=1
           VCCV Control Verification types:
             LSP ping
             BFD with PW-ACH-encapsulation for Fault Detection
             BFD with IP/UDP-encapsulation for Fault Detection

Output label database, 10.255.107.232:0--10.255.107.236:0
  Label Prefix
  3 10.255.107.232/32
   State: Active
   Age: 9:35
   Entropy Label Capability: No
  299776 10.255.107.236/32

```

State: Active
 Age: 9:35
 Entropy Label Capability: No

show ldp database (standby)

user@host> show ldp database extensive

```
Input label database, 10.255.107.236:0--10.255.107.234:0
Label Prefix
299808 10.255.107.230/32
      State: Active
      Age: 1d 2:46:36
      Standby binding state:
        Map messages: 1
        Release messages: 0
Label Prefix
301136 10.255.107.232/32
      State: Active
      Age: 1d 2:46:36
      Standby binding state:
        Map messages: 1
        Release messages: 0
Label Prefix
3      10.255.107.234/32
      State: Active
      Age: 1d 2:46:36
      Standby binding state:
        Map messages: 1
        Release messages: 0
Label Prefix
302480 10.255.107.236/32
      State: Active
      Age: 1d 2:46:36
      Standby binding state:
        Map messages: 1
        Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.234:0
Label Prefix
299904 10.255.107.230/32
      State: Active
      Age: 1d 2:46:36
299936 10.255.107.232/32
      State: Active
      Age: 1d 2:46:36
299872 10.255.107.234/32
      State: Active
      Age: 1d 2:46:36
3      10.255.107.236/32
      State: Active
      Age: 1d 2:46:36
299952 P2MP root-addr 10.255.107.230, lsp-id 16777217
      State: Active
      Age: 1d 2:46:36
```

show ldp database l2circuit detail

user@host> show ldp database l2circuit detail

```
Input label database, 10.255.245.44:0--10.255.245.45:0
Label Prefix
```



```

100176      L2CKT CtrlWord ATM CELL (VC Mode) VC 100
           Cell bundle size: 80
           State: Active
           Age: 9:48
100256      L2CKT CtrlWord FRAME RELAY VC 101
           MTU: 4470
           State: Active
           Age: 9:48

Output label database, 10.255.245.44:0--10.255.245.45:0
Label      Prefix
100048      L2CKT CtrlWord ATM CELL (VC Mode) VC 100
           Cell bundle size: 80
           State: Active
           Age: 9:48
100112      L2CKT CtrlWord FRAME RELAY VC 101
           MTU: 4470
           State: Active
           Age: 9:48

```

show ldp database l2circuit extensive

```

user@host> show ldp database l2circuit extensive
Input label database, 10.255.245.198:0--10.255.245.194:0
Label      Prefix
299872      L2CKT CtrlWord PPP VC 100
           MTU: 4470
           VCCV Control Channel types:
             MPLS router alert label
             MPLS PW label with TTL=1
           VCCV Control Verification types:
             LSP ping
Label      Prefix
           State: Active
           Age: 19:23:08

```

show ldp database p2mp (master)

```

user@host> show ldp database p2mp extensive

Input label database, 10.255.107.232:0--10.255.107.236:0
Label      Prefix
569649      P2MP root-addr 10.255.107.232, lsp-id 16777217
           State: Active
           Age: 2d 6:41:46

Output label database, 10.255.107.232:0--10.255.107.236:0

Input label database, 10.255.107.232:0--10.255.107.238:0

Output label database, 10.255.107.232:0--10.255.107.238:0
Label      Prefix
299888      P2MP root-addr 10.255.107.230, lsp-id 16777217
           State: Active
           Age: 2d 6:41:35

```

show ldp database p2mp (standby)

```

user@host> show ldp database p2mp extensive

Input label database, 10.255.107.236:0--10.255.107.232:0

```

```

Label      Prefix
299968     P2MP root-addr 10.255.107.230, lsp-id 16777217
           State: Active
           Age: 4d 22:21:57
           Standby binding state:
             Map messages: 1
             Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.232:0
Label      Prefix
3          P2MP root-addr 10.255.107.232, lsp-id 1
           State: Active
           Age: 4d 22:21:57

```

show ldp database p2mp (master)

```
user@host> show ldp database p2mp extensive
```

```

Input label database, 10.255.107.232:0--10.255.107.236:0
Label      Prefix
569649     P2MP root-addr 10.255.107.232, lsp-id 16777217
           State: Active
           Age: 2d 6:41:46

Output label database, 10.255.107.232:0--10.255.107.236:0

Input label database, 10.255.107.232:0--10.255.107.238:0

Output label database, 10.255.107.232:0--10.255.107.238:0
Label      Prefix
299888     P2MP root-addr 10.255.107.230, lsp-id 16777217
           State: Active
           Age: 2d 6:41:35

```

show ldp database p2mp (standby)

```
user@host> show ldp database p2mp extensive
```

```

Input label database, 10.255.107.236:0--10.255.107.232:0
Label      Prefix
299968     P2MP root-addr 10.255.107.230, lsp-id 16777217
           State: Active
           Age: 4d 22:21:57
           Standby binding state:
             Map messages: 1
             Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.232:0
Label      Prefix
3          P2MP root-addr 10.255.107.232, lsp-id 1
           State: Active
           Age: 4d 22:21:57

```

show ldp database session

```

user@host> show ldp database session 10.1.1.195
Input label database, 10.0.0.194:0--10.1.1.195:0
Label      Prefix
100002     10.255.245.197/32
100003     10.255.245.196/32
100004     10.0.0.194/32

```

```

      3      10.1.1.195/32
100000      L2CKT NoCtrlWord FRAME RELAY VC 1
100001      L2CKT CtrlWord FRAME RELAY VC 2
Output label database, 10.0.0.194:0--10.1.1.195:0
  Label      Prefix
100003      10.255.245.197/32
100004      10.1.1.195/32
100002      10.255.245.196/32
      3      10.0.0.194/32
100000      L2CKT CtrlWord FRAME RELAY VC 2
100001      L2CKT NoCtrlWord FRAME RELAY VC 1

```

show ldp database (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show ldp database
Input label database, 1.1.1.2:0--1.1.1.3:0
  Label      Prefix
299808      1.1.1.2/32
      3      1.1.1.3/32
299792      1.1.1.6/32
299776      10.255.2.227/32
299840      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
299824      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11

Output label database, 1.1.1.2:0--1.1.1.3:0
  Label      Prefix
      3      1.1.1.2/32
299776      1.1.1.3/32
299808      1.1.1.6/32
299792      10.255.2.227/32

Input label database, 1.1.1.2:0--1.1.1.6:0
  Label      Prefix
299856      1.1.1.2/32
299792      1.1.1.3/32
      3      1.1.1.6/32
299776      10.255.2.227/32
299888      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
299808      P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
299824      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
299840      P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
299872      P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7

Output label database, 1.1.1.2:0--1.1.1.6:0
  Label      Prefix
      3      1.1.1.2/32
299776      1.1.1.3/32
299808      1.1.1.6/32
299792      10.255.2.227/32

```

show ldp database (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show ldp database
Input label database, 10.255.2.227:0--1.1.1.3:0
  Label      Prefix
299808      1.1.1.2/32
      3      1.1.1.3/32
299792      1.1.1.6/32
299776      10.255.2.227/32

Output label database, 10.255.2.227:0--1.1.1.3:0
  Label      Prefix

```

```

299856      1.1.1.2/32
299776      1.1.1.3/32
299792      1.1.1.6/32
3           10.255.2.227/32

```

Input label database, 10.255.2.227:0--1.1.1.6:0

```

Label      Prefix
299856      1.1.1.2/32
299776      1.1.1.3/32
3           1.1.1.6/32
299776      10.255.2.227/32

```

Output label database, 10.255.2.227:0--1.1.1.6:0

```

Label      Prefix
299856      1.1.1.2/32
299776      1.1.1.3/32
299792      1.1.1.6/32
3           10.255.2.227/32
299888      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
299808      P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
299824      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
299840      P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
299872      P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7

```

show ldp database summary

```
user@host> show ldp database summary
```

Session ID	Labels received	Labels advertised
10.255.0.1:0--10.255.0.2:0	4	4
10.255.0.1:0--10.255.0.3:0	4	4

show ldp fec-filters

Syntax	show ldp fec-filters <fec> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Display information about configured Label Distribution Protocol (LDP) forwarding equivalence class (FEC) filters.
Options	<p>fec—(Optional) Display FEC filter information for the specified FEC.</p> <p>instance <i>instance-name</i>—(Optional) Display FEC filter information for the specified instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show ldp fec-filters on page 1319
Output Fields	Table 56 on page 1319 lists the output fields for the show ldp fec-filters command. Output fields are listed in the approximate order in which they appear.

Table 56: show ldp fec-filters Output Fields

Field Name	Field Description
Ingress	Names of the FEC filters on the ingress routers.
Transit	Names of the FEC filters on the transit routers.

Sample Output

show ldp fec-filters

```

user@host> show ldp fec-filters 10/8
10.22.1.2/32
  Ingress: f1-10.22.1.2/32 (index: 3)
  Transit: (null) (index: 0)

```

show ldp interface

Syntax	<pre>show ldp interface <brief detail extensive> <interface-name> <instance instance-name> <logical-system (all logical-system-name)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	Display the status of Label Distribution Protocol (LDP)-enabled interfaces.
Options	<p>none—Display standard status information about all LDP-enabled interface for all routing instances.</p> <p>interface-name—(Optional) Display information for the specified interface.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance instance-name—(Optional) Display information for the specified routing instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show ldp interface extensive on page 1321
Output Fields	<p>Table 57 on page 1320 describes the output fields for the show ldp interface command. Output fields are listed in the approximate order in which they appear.</p>

Table 57: show ldp interface Output Fields

Field Name	Field Description	Level of Output
Interface	Interface name.	All levels
Label space ID	Label space identifier that the router is advertising on the interface.	All levels
Nbr count	Number of neighbors on the interface.	All levels
Next hello	How long until the next hello packet is sent on this interface, in seconds.	All levels
Hello interval	One-third of the negotiated hold time (in seconds). If the user-configured value for the hello interval is smaller than the computed value, the user-configured value is used.	detail extensive
Hold time	Configured hold time, in seconds.	detail extensive

Table 57: show ldp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transport address	Address to which the neighbor wants the local route to establish the LDP session.	extensive
Local hello interval	Locally configured hello interval.	extensive

Sample Output

show ldp interface extensive

```
user@host> show ldp interface extensive
Interface          Label space ID      Nbr count  Next hello
fe-0/0/3.0         10.255.245.6:0      2          0
Hello interval: 1, Hold time: 15, Transport address: 10.255.245.6
Local hello interval: 2, Index: 69
```

show ldp neighbor

Syntax	<pre>show ldp neighbor <brief detail extensive> <auto-targeted> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor-address></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>neighbor-address option added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p> <p>auto-targeted option added in Junos OS Release 14.2.</p>
Description	Display Label Distribution Protocol (LDP) neighbor information.
Options	<p>none—Display standard information about LDP neighbors for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>auto-targeted—(Optional) Display information about LDP neighbors that are automatically targeted using the loopback addresses.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor-address—(Optional) Display information about the specified LDP neighbor.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear ldp neighbor on page 1304
List of Sample Output	<p>show ldp neighbor extensive on page 1323</p> <p>show ldp neighbor auto-targeted extensive on page 1323</p>
Output Fields	<p>Table 58 on page 1322 describes the output fields for the show ldp neighbor command. Output fields are listed in the approximate order in which they appear.</p>

Table 58: show ldp neighbor Output Fields

Field Name	Field Description	Level of Output
Address	IP address of the neighbor.	All levels
Interface	Interface over which the neighbor was discovered.	All levels
Label space ID	Label space identifier advertised by the neighbor.	All levels

Table 58: show ldp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Hold time	Remaining hold time before the neighbor expires, in seconds.	All levels
Transport address	Address to which the neighbor wants the local route to establish the LDP session.	detail
Configuration sequence	Counter that increments whenever the neighbor changes its configuration.	detail
Up for	Length of time the LDP neighbor has been in operation.	detail extensive
Reference count	Reference count for the LDP neighbor.	extensive
Hold time	Displays the neighbor's hold time. The hold time is the proposed hold times for the local and peer routers.	extensive
Proposed local/peer	Hold time value proposed by the local router and the peer router.	extensive

Sample Output

show ldp neighbor extensive

```

user@host> show ldp neighbor extensive
Address          Interface      Label space ID      Hold Time
192.168.37.23    so-1/0/0.0    10.255.245.5:0      44
Transport address: 10.255.245.5, Configuration sequence: 6
Up for 00:03:37
Reference count: 1
Hold time: 45, Proposed local/peer: 15/45

```

show ldp neighbor auto-targeted extensive

```

user@host> show ldp neighbor auto-targeted extensive
Address          Interface      Label space ID      Hold time
10.255.107.236   lo0.0         10.255.107.236:0    41
Transport address: 10.255.107.236, Configuration sequence: 14
Up for 00:10:53
Reference count: 2
Hold time: 45, Proposed local/peer: 45/45
Hello interval: 15
Hello flags: targeted
Neighbor types: Auto-targeted

```

show ldp overview

Syntax	show ldp overview <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	show ldp overview <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 11.2.
Description	Display LDP overview information.
Options	<p>none— Display standard overview information about LDP for all routing instances.</p> <p>instance <i>instance-name</i>— (Optional) Display LDP overview information for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)— (Optional) Display LDP information from systems or a particular logical system on the devices.</p>
Required Privilege Level	view
List of Sample Output	show ldp overview on page 1327
Output Fields	Table 59 on page 1324 lists the output fields for the show ldp overview command. Output fields are listed in the approximate order in which they appear.

Table 59: show ldp overview Output Fields

Field Name	Field description	Level of Output
Instance	LDP routing instance.	All Levels
Router ID	Router ID of the routing device.	All Levels
Message ID	Unique identifier of message.	All Levels
Configuration sequence	Value of configuration sequence.	All Levels
Deaggregate	Status of control forwarding equivalence class (FEC) deaggregation on the router. By default it is disabled on the router.	All Levels
Explicit null	<p>Advertise label 0 to the egress routing device of an LSP. Explicit null: enabled or disabled.</p> <p>NOTE: If you do not include the explicit-null statement in the configuration, label 3 (implicit null) is advertised.</p>	All Levels
IPv6 tunneling	Internet Protocol version 6 tunneling: enabled or disabled .	All Levels

Table 59: show ldp overview Output Fields (*continued*)

Field Name	Field description	Level of Output
Strict targeted hellos	Prevent LDP sessions from being established with remote neighbors that have not been specifically configured. Strict targeted hellos: enabled or disabled . <i>NOTE:</i> LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.	All Levels
Loopback if added	Loopback interface is added: yes or no .	All Levels
Route preference	Default preference value (also known as an administrative distance) assigned to each route that the routing table receives. LDP preference is: 9	All Levels
Unicast transit LSP chaining	Unicast transit LSP chaining: enabled or disabled .	All Levels
P2MP transit LSP chaining	P2MP transit LSP chaining: enabled or disabled .	All Levels
Transit LSP statistics based on route statistics	Transit LSP statistics based on route statistics: enabled or disabled .	All Levels
Capabilities enabled	Enabled capabilities: none	All Levels
Timers	<ul style="list-style-type: none"> • Keepalive interval: Keepalive interval value. • Keepalive timeout: Time interval for which the neighbor LDP node waits before determining session failure. • Link hello interval: Specify how often the router sends Link Management Protocol (LMP) hello packets. • Link hello hold time: Time interval for which an LDP node waits for a hello message before declaring a neighbor is down. • Targeted hello interval: Specify how often LDP sends targeted hello messages. • Targeted hello hold time: Time interval for which a sending LSR maintains a record of targeted hello messages from the receiving LSR without receipt of another targeted hello message from that LSR. • Label withdraw delay: Time interval for withdrawing labels to reduce router workload during IGP convergence. 	All Levels

Table 59: show ldp overview Output Fields (*continued*)

Field Name	Field description	Level of Output
Graceful restart	Graceful restart attributes. <ul style="list-style-type: none"> • Restart— Graceful restart capability: enabled or disabled. • Helper— Standard graceful restart helper capability: enabled or disabled. • Restart in process— Graceful restart in process. • Reconnect time— Period of time that a restarting LSR (label switched router) designates to LDP neighbors to wait until the former reestablishes the session after restarting. • Max neighbor reconnect time— Maximum reconnect time. • Recovery time— Period of time that an LSR preserves its state across the restart. • Max neighbor recovery time— Maximum recovery time designated to LDP neighbors by a restarting LSR. 	All Levels
Traffic Engineering	<ul style="list-style-type: none"> • Bgp igp— BGP and IGP destinations: enabled or disabled. When enabled, IGP uses MPLS paths for forwarding traffic. • Both ribs— BGP and IGP destinations with routes in both RIBs: enabled or disabled. • Mpls forwarding— MPLS routes used for forwarding: enabled or disabled. 	All Levels
IGP	<ul style="list-style-type: none"> • Tracking igp metric— Cause the IGP route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1). • Sync session up delay— Time interval to synchronize LDP session. 	All Levels
Session protection	<ul style="list-style-type: none"> • Session protection— Remote neighbor added to LDP configuration which enables protection for all sessions in the corresponding LDP instance: enabled or disabled. • Session protection timeout— Period of time until which the remote neighbor is connected to LSR in the absence of link neighbors. 	All Levels
Interface addresses advertising	Advertises interface address.	All Levels
Label allocation	Label accounting information. <ul style="list-style-type: none"> • Current number of LDP labels allocated— Number of labels currently in use. • Total number of LDP labels allocated— Cumulative number of labels being allocated. • Total number of LDP labels freed— Cumulative number of labels being freed. • Total number of LDP label allocation failure— Cumulative number of failures for allocating a label • Current number of labels allocated by all protocols— Number of labels currently being used by routing protocols. 	All Levels

Sample Output

show ldp overview

```
user@host> show ldp overview
```

Sample Output

```
Instance: master
Router ID: 192.168.2.1
Message id: 0
Configuration sequence: 1
Deaggregate: disabled
Explicit null: disabled
IPv6 tunneling: disabled
Strict targeted hellos: disabled
Loopback if added: yes
Route preference: 9
Unicast transit LSP chaining: disabled
P2MP transit LSP chaining: disabled
Transit LSP statistics based on route statistics: disabled
Capabilities enabled: none
Timers:
  Keepalive interval: 10, Keepalive timeout: 30
  Link hello interval: 5, Link hello hold time: 15
  Targeted hello interval: 15, Targeted hello hold time: 45
  Label withdraw delay: 60
Graceful restart:
  Restart: enabled, Helper: enabled, Restart in process: false
  Reconnect time: 60000, Max neighbor reconnect time: 120000
  Recovery time: 160000, Max neighbor recovery time: 240000
Traffic Engineering:
  Bgp igp: disabled
  Both ribs: disabled
  Mpls forwarding: disabled
IGP:
  Tracking igp metric: disabled
  Sync session up delay: 10
Session protection:
  Session protection: disabled
  Session protecton timeout: 0
Interface addresses advertising:
  192.168.2.1
Label allocation:
  Current number of LDP labels allocated: 3
  Total number of LDP labels allocated: 3
  Total number of LDP labels freed: 0
  Total number of LDP label allocation failure: 0
  Current number of labels allocated by all protocols: 3
```

show ldp p2mp tunnel

Syntax	show ldp p2mp tunnel <brief detail extensive> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced in Junos OS Release 13.3.
Description	Display LDP point-to-multipoint tunnel table information.
Options	brief detail extensive —(Optional) Display the specified level of output. instance <i>instance-name</i> —(Optional) Display routing instance information for the specified instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Display LDP point-to-multipoint tunnel table information of all logical systems or a particular logical system.
Required Privilege Level	View

Sample Output

show ldp p2mp tunnel

```
user@host> show ldp p2mp tunnel extensive
```

```
Instance      Tunnel type      Tunnel name
0             Name            10.254.1.1:1:ldp-p2mp:mvpn:vpn-1
P2MP root-addr 10.255.107.232, lsp-id 16777217
Self id 805306372
Reference count 2
```

show ldp path

Syntax	show ldp path <brief detail extensive> <destination> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Display Label Distribution Protocol (LDP) label-switched paths (LSPs).
Options	<p>none—Display standard information about all LDP LSPs for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>destination—(Optional) Restrict the output to entries that match the specified destination prefix.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show ldp path extensive on page 1330
Output Fields	Table 60 on page 1329 describes the output fields for the show ldp path command. Output fields are listed in the approximate order in which they appear.

Table 60: show ldp path Output Fields

Field Name	Field Description
Output Session (label)	Session ID and labels that this system has sent using LDP. These correspond to MPLS packets received.
Input Session (label)	Session ID and labels that this system has received using LDP. These correspond to MPLS packets transmitted.
route	MPLS route.
Attached route	Route corresponding to the LSP.
Ingress route	The router acts as the ingress for the LSP.
Reference count	Reference count for the LDP neighbor.

Table 60: show ldp path Output Fields (*continued*)

Field Name	Field Description
Transit route	Names of the forwarding equivalence class (FEC) filters on the transit routers.
Global label	MPLS label that is used globally.

Sample Output

show ldp path extensive

```
user@host> show ldp path extensive
Output Session (label)      Input Session (label)
10.255.14.220:0(3)         ( )
  Attached route: 10.255.14.221/32
  Reference count: 3, Global label: 3
10.255.14.220:0(100000)     10.255.14.220:0(3)
  Attached route: 10.255.14.220/32, Ingress route
  Reference count: 2, Transit route, Global label: 100000
10.255.14.220:0(100001)     10.255.14.220:0(100001)
  Attached route: 10.255.14.214/32, Ingress route
  Reference count: 2, Transit route, Global label: 100001
```


show ldp route

Syntax	<pre>show ldp route <brief detail extensive> <destination> <instance instance-name> <logical-system (all logical-system-name)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
Description	Display the entries in the Label Distribution Protocol (LDP) internal topology table. The internal topology table contains routes from inet.0 and inet.3 and is used when binding a label to a forwarding equivalence class (FEC).
Options	<p>none—Display standard information about all entries in the LDP internal topology table for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>destination—(Optional) Restrict the output to entries that are longer than the specified destination prefix and prefix length.</p> <p>instance instance-name—(Optional) Display entries for the specified routing instance only.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	<p>show ldp route detail on page 1333</p> <p>show ldp route extensive on page 1333</p>
Output Fields	Table 61 on page 1331 describes the output fields for the show ldp route command. Output fields are listed in the approximate order in which they appear.

Table 61: show ldp route Output Fields

Field Name	Field Description
Destination	Destination prefix.
Next-hop intf/lsp/table	Interface that is the next hop to the destination prefix.
Next-hop address	IP address of the next hop.
Session ID	LDP session ID.

Table 61: show ldp route Output Fields (*continued*)

Field Name	Field Description
Route flags	Information about the route. For example, the Ingress TTL propagate flag indicates that the time-to-live (TTL) value is being propagated with the route.
Bound to outgoing label	The route has been bound to LSPs with the label being distributed for that LSP.
Topology entry	The topology that the route is bound to.
Ingress route status	Status of the ingress route. For example, it could be Active or Inactive .
Last modified	The length of time since the ingress route status last changed.

Sample Output

show ldp route detail

```

user@host> show ldp route 10.255.8.5 detail
Destination      Next-hop intf/lsp      Next-hop address
10.255.8.5/32     f1
  Session ID 10.255.170.84:0--10.255.170.92:0
                    fe-0/0/0.0      192.168.100.2
  Session ID 10.255.170.84:0--10.255.8.5:0
                    so-0/2/1.0
  Session ID 10.255.170.84:0--10.255.8.5:0
                    so-0/2/2.0
  Session ID 10.255.170.84:0--10.255.8.3:0
  Bound to outgoing label 299776, Topology entry: 0x8c38a80
  BFD dest addr   BFD state LSP-ping Next-hop addr Next-hop intf/lsp
127.0.0.64       up        up        192.168.100.2 fe-0/0/0.0
127.0.1.64       up        up        so-0/2/1.0
127.0.2.64       up        up        so-0/2/2.0
127.0.3.64       up        up        f1
.....

```

show ldp route extensive

```

user@host> show ldp route extensive

Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.0/30      ge-1/2/0.18            10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.4/30      ge-1/2/0.18            10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.8/30      ge-1/2/1.21            10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.12/30     ge-1/2/1.21            10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.16/30     ge-1/2/0.18
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.18/32     ge-1/2/0.18
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.20/30     ge-1/2/1.21
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.21/32     ge-1/2/1.21
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
192.168.0.1/32   ge-1/2/0.18            10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
192.168.0.2/32   ge-1/2/1.21            10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0

```

```

                                ge-1/2/0.18                10.0.0.17
    Session ID 192.168.0.6:0--192.168.0.5:0
    Route flags: None
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.3/32   ge-1/2/1.21                10.0.0.22
    Session ID 192.168.0.6:0--192.168.0.4:0
    Route flags: None
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.4/32   ge-1/2/1.21                10.0.0.22
    Session ID 192.168.0.6:0--192.168.0.4:0
    Bound to outgoing label 299808, Topology entry: 0x92a483c
    Ingress route status: Active, Last modified: 00:01:19 ago
    Route flags: Ingress TTL propagate, Transit TTL propagate
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.5/32   ge-1/2/0.18                10.0.0.17
    Session ID 192.168.0.6:0--192.168.0.5:0
    Bound to outgoing label 299792, Topology entry: 0x92a47f8
    Ingress route status: Active, Last modified: 00:01:19 ago
    Route flags: Ingress TTL propagate, Transit TTL propagate
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.6/32   lo0.6
    Bound to outgoing label 3, Topology entry: 0x92a4a5c
    Ingress route status: Inactive
    Route type: Egress route
    Route flags: None
Destination      Next-hop intf/lsp/table      Next-hop address
10.10.20.1/32    fe-1/0/0.0                192.168.199.37
                                LSP LDP->10.255.107.230

```

show ldp session

Syntax	<pre>show ldp session <brief detail extensive> <auto-targeted> <destination> <instance instance-name> <logical-system (all logical-system-name)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p> <p>auto-targeted option added in Junos OS Release 14.2.</p>
Description	Display information about Label Distribution Protocol (LDP) sessions.
Options	<p>none—Display standard information about all LDP sessions for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>auto-targeted—(Optional) Display information about LDP sessions that are automatically targeted using loopback addresses.</p> <p>destination—(Optional) Restrict LDP session display to the specified address.</p> <p>instance instance-name—(Optional) Display routing instance information for the specified instance. If instance-name is omitted, information is displayed for the master instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ldp session on page 1305
List of Sample Output	<p>show ldp session brief on page 1339</p> <p>show ldp session detail on page 1339</p> <p>show ldp session extensive on page 1339</p> <p>show ldp session auto-targeted detail on page 1340</p>
Output Fields	<p>Table 62 on page 1335 describes the output fields for the show ldp session command. Output fields are listed in the approximate order in which they appear.</p>

Table 62: show ldp session Output Fields

Field Name	Field Description	Level of Output
Address	Transport address of the session.	any
State	State of the session: Nonexistent , Connecting , Initialized , OpenRec , OpenSent , Operational , or Closing . The states correspond to the state diagram specified in Internet Draft LDP Specification draft-ietf-mpls-rfc3036bis-01.txt.	any

Table 62: show ldp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Connection	TCP connection state: Closed , Opening , or Open .	any
Hold time	Time remaining until the session will be closed, in seconds.	any
Session ID	LDP identifiers of the peers of this session.	detail extensive
Next keepalive	Time until next keepalive is sent, in seconds.	detail extensive
Active	Whether the local router is playing the active role in the session and during session establishment.	detail extensive
Passive	Whether the local router is playing the passive role in the session and during session establishment.	detail extensive
Maximum PDU	Maximum protocol data unit (PDU) size (packet size) for the session.	detail extensive
Hold time	Time remaining until the session will be closed, in seconds. This value corresponds to the one configured using the keepalive-timeout statement configured at the [edit protocols ldp] hierarchy level.	detail extensive
Neighbor count	Number of neighbors that are contributing to the session.	detail extensive
Neighbor types	Category of LDP session: discovered or auto-targeted .	any
Keepalive interval	Keepalive interval, in seconds.	detail extensive
Connect retry interval	TCP connection retry interval, in seconds.	detail extensive
Local address	Local transport address.	detail extensive
Remote address	Remote transport address.	detail extensive
Up for	Time that this session has been up.	detail extensive
Last down	Time since the session last went down.	detail extensive

Table 62: show ldp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Reason	Reason the session went down: <ul style="list-style-type: none"> • Aborted graceful restart • Authentication key was changed • Bad type length value (TLV) • Bad protocol data unit (PDU) packets • Command-line interface (CLI) command • Connect time expired • Connection error • Connection reset • Error during initialization • Hold time expired • No adjacency or all adjacencies down • Notification received • Received notification from peer • Unexpected End of File (EOF) • Unknown reason 	detail extensive
Number of session flaps	Number of times the session changes from up to down.	detail extensive
Restarting	LDP is in the process of gracefully restarting.	detail extensive
Capabilities advertised	LDP capabilities advertised to a peer.	detail extensive
Capabilities received	LDP capabilities received from a peer.	detail extensive
Protection	Information about the status of MPLS LDP session protection.	detail extensive
restart complete in <i>nnn msec</i>	Amount of time (in milliseconds) remaining until graceful restart is declared complete.	detail extensive
Local	Information about graceful restart for the local end of an LDP session. Graceful restart and helper mode are independent. <ul style="list-style-type: none"> • Restart—Status of the graceful restart feature at the local end of the LDP session: enabled or disabled. • Helper mode—Status of the helper mode feature at the local end of the LDP session: enabled or disabled. When this feature is enabled, the local end of the LDP session can help the restarting router with its LDP restart procedures. • Reconnect time—Amount of time to wait from when a restart is initiated until the router can exchange LDP messages with its neighbors. The default is 60000 msec and is not configurable. (Reconnect timeout refers to "FT Reconnect timeout" in draft-ietf-mpls-ldp-restart-06, <i>Internet Draft Graceful Restart Mechanism for LDP</i>.) 	detail extensive

Table 62: show ldp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Remote	<p>Information about graceful restart at the remote end of an LDP session. Graceful restart and helper mode are independent.</p> <ul style="list-style-type: none"> • Restart—Status of the graceful restart feature at the remote end of the LDP session: enabled or disabled. • Helper mode—Status of the helper mode feature at the remote end of the LDP session: enabled or disabled. When this feature is enabled, the remote end of the LDP session can help the restarting router with its LDP restart procedures. • Reconnect time—Amount of time in milliseconds from when a restart is initiated until the remote router can exchange LDP messages with its neighbors. 	detail extensive
Local maximum recovery time	Amount of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds).	detail extensive
Next-hop addresses received	Next-hop addresses received on the session.	detail extensive
Queue depth	Number of messages that are queued for sending to the peers in the group.	extensive
Message type	<p>Type of message being sent:</p> <ul style="list-style-type: none"> • Initialization—Session initialization negotiation messages sent by an LSR to an LDP peer when the transport connection is established. • Keepalive—Keepalive timer messages sent by an LSR to an LDP peer to keep the session active when there is no information or PDU exchanged between them. • Notification—Notification messages (such as state of the LDP session) or error information (such as bad PDU length) sent by an LSR to an LDP peer. • Address—Message sent by an LSR to an LDP peer to advertise interface addresses. • Address withdraw—Message sent by an LSR to an LDP peer to withdraw a previously advertised interface address. • Label mapping—Message sent by an LSR to an LDP peer to advertise label mapping for a forwarding equivalence class (FEC). • Label request—Message sent by an LSR to an LDP peer to request a label mapping for an FEC. • Label withdraw—Message sent by an LSR to an LDP peer to withdraw a previously advertised FEC-label mapping. • Label release—Message sent by an LSR to an LDP peer to notify the peer that a specific FEC-label mapping has been released. • Label abort—Message sent by an LSR to an LDP peer to abort a label request message. • Total—Messages sent and received during the lifetime of the session. • Last 5 seconds—Messages sent and received during the current session. 	extensive

Sample Output

show ldp session brief

```
user@host> show ldp session brief
  Address           State           Connection      Hold time
10.255.72.160       Operational     Open            21
10.255.72.164       Operational     Open            20
10.255.72.172       Operational     Open            21
```

show ldp session detail

```
user@host> show ldp session detail
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 27
Session ID: 192.168.0.2:0--192.168.0.3:0
Next keepalive in 7 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.0.2, Remote address: 192.168.0.3
Up for 00:00:02
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
  10.0.0.5
  10.0.0.33
```

show ldp session extensive

```
user@host> show ldp session extensive
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 22
Session ID: 192.168.0.2:0--192.168.0.3:0
Next keepalive in 2 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.0.2, Remote address: 192.168.0.3
Up for 00:05:37
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
  10.0.0.5
  10.0.0.33
```

Queue depth: 0				
Message type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Initialization	1	1	0	0
Keepalive	33	33	1	1
Notification	0	0	0	0
Address	1	1	0	0
Address withdraw	0	0	0	0
Label mapping	7	5	0	0
Label request	0	0	0	0
Label withdraw	3	1	0	0
Label release	1	3	0	0
Label abort	0	0	0	0

show ldp session auto-targeted detail

```

user@host> show ldp session auto-generated detail
Address: 192.168.1.5, State: Operational, Connection: Open, Hold time: 25
  Session ID: 192.168.1.1:0--192.168.1.5:0
  Next keepalive in 5 seconds
  Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Neighbor types: discovered, Auto-targeted
                    ^^^^^^^^^^^^^^^^^
  Keepalive interval: 10, Connect retry interval: 1
  Local address: 192.168.1.1, Remote address: 192.168.1.5
  Up for 00:00:34
  Capabilities advertised: none
  Capabilities received: none
  Protection: disabled
  Local - Restart: disabled, Helper mode: enabled
  Remote - Restart: disabled, Helper mode: enabled
  Local maximum neighbor reconnect time: 120000 msec
  Local maximum neighbor recovery time: 240000 msec
  Local Label Advertisement mode: Downstream unsolicited
  Remote Label Advertisement mode: Downstream unsolicited
  Negotiated Label Advertisement mode: Downstream unsolicited
  Nonstop routing state: Not in sync
  Next-hop addresses received:
    192.168.1.2
    192.168.1.3

```

show ldp statistics

Syntax	show ldp statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Display Label Distribution Protocol (LDP) statistics.
Options	<p>none—Display LDP statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ldp statistics on page 1306
List of Sample Output	show ldp statistics on page 1344
Output Fields	Table 63 on page 1341 lists the output fields for the show ldp statistics command. Output fields are listed in the approximate order in which they appear.

Table 63: show ldp statistics Output Fields

Field Name	Field Description
Total Sent, Received	Total number of each message type sent and received.
Last 5 seconds Sent, Received	Number of each message type sent and received in the last 5 seconds.

Table 63: show ldp statistics Output Fields (*continued*)

Field Name	Field Description
Message type	<p>LDP message types:</p> <ul style="list-style-type: none"> • Hello—Messages that enable LDP nodes to discover one another and to detect the failure of a neighbor or of the link to the neighbor. • Initialization—Messages that indicate an LDP session has started. • Keepalive—Messages that ensure that the keepalive timeout is not exceeded. • Notification—Advisory information and signal error information. • Address—Messages with address information. • Address withdrawal—Messages regarding address withdrawal. • Label mapping—Messages with label mapping information. • Label request—Request for a label mapping from a neighboring router. • Label withdrawal—Withdrawal message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use. • Label release—Message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use. • Label abort—Messages about label interruptions. • All UDP—All hello messages sent by LSRs to the well-known UDP port, 646. • All TCP—All LDP session messages.

Table 63: show ldp statistics Output Fields (*continued*)

Field Name	Field Description
Event type	<p>LDP events and errors:</p> <ul style="list-style-type: none"> • Sessions opened—Number of LDP sessions that have been opened. • Sessions closed—Number of LDP sessions that have been closed. • Topology changes—Number of changes to the known LDP topology. • No interface—Number of missing interface address messages. When a new LDP session is initialized and before sending label lapping or label request messages, the LSR advertises its interface addresses with one or more address messages. • No session—Number of missing session messages. Session messages are used to establish, maintain, and terminate sessions between LDP peers. • No adjacency—The exchange of hello adjacency messages results in the creation of an adjacency. The LDP identifier, together with the sender's LDP identifier in the PDU header, enables the receiver to match the initialization message with one of its hello adjacencies. If there is no matching hello adjacency, the LSR sends a session the initialization message is rejected. • Unknown version—The LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment. • Malformed PDU—An LDP PDU received on a TCP connection for an LDP session is malformed if the LDP identifier in the PDU header is unknown to the receiver, or if it is known but is not the LDP identifier associated by the receiver with the LDP peer for this LDP session. An LDP PDU is considered to be malformed if the LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment. An LDP PDU is considered malformed if the PDU length field is too small (less than 14) or too large (greater than maximum PDU length). • Malformed message—Malformed LDP messages that are part of the LDP discovery mechanism are handled by silently discarding them. An LDP message is malformed if the message type is unknown. If the message type is less than 0x8000 (high order bit = 0), it is an error signaled by the unknown message type status code. An LDP message is considered to be malformed if the message length is too large, meaning that the message extends beyond the end of the containing LDP PDU. The LDP message is considered to be malformed if the message length is too small, meaning that it is smaller than the smallest possible value component. The LDP message is considered to be malformed if the message is missing one or more mandatory parameters. • Unknown message type—If the message type is less than 0x8000 (high order bit = 0) or greater than or equal to 0x8000 (high order bit = 1) it is considered to be an unknown message. • Inappropriate message—The message is not of the type that the receiver expects to receive. • Malformed TLV—The TLV Length is too large or the receiver cannot decode the TLV value. This can indicate an issue in either the sending or receiving LSR. • Bad TLV value—The TLV Length is too large. • Missing TLV—The TLV is missing one or more mandatory parameters. • PDU too large—The PDF is greater than the maximum PDU length. Section "Initialization Message" in RFC 5036 describes how the maximum PDU length for a session is determined.
Total	Total number of each event or error.
Last 5 seconds	Number of each event or error in the last 5 seconds.

Sample Output


show ldp statistics

```
user@host> show ldp statistics
```

Message type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Hello	265	263	2	2
Initialization	2	2	0	0
Keepalive	112	111	1	0
Notification	0	0	0	0
Address	2	2	0	0
Address withdraw	0	0	0	0
Label mapping	7	6	0	0
Label request	0	0	0	0
Label withdraw	2	0	0	0
Label release	0	2	0	0
Label abort	0	0	0	0
All UDP	265	263	2	2
All TCP	123	121	1	0

Event type	Total	Last 5 seconds	
		Sent	Received
Sessions opened	2		0
Sessions closed	0		0
Topology changes	11		0
No interface	0		0
No session	0		0
No adjacency	0		0
Unknown version	0		0
Malformed PDU	0		0
Malformed message	0		0
Unknown message type	0		0
Inappropriate message	0		0
Malformed TLV	0		0
Bad TLV value	0		0
Missing TLV	0		0
PDU too large	0		0

show ldp traffic-statistics

Syntax	<pre>show ldp traffic-statistics <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <p2mp></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>p2mp option added in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p>
Description	Display Label Distribution Protocol (LDP) traffic statistics.
<div>  <p>NOTE: If nonstop active routing features is configured, show ldp traffic-statistics command is not supported on backup Routing Engines.</p> </div>	
Options	<p>none—Display LDP traffic statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display LDP traffic statistics for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>p2mp—(Optional) Display only the data traffic statistics for a point-to-multipoint LSP.</p>
Additional Information	To collect output from this command on a periodic basis, configure the traffic-statistics statement for the LDP protocol. For more information, see the <i>Junos MPLS Applications Configuration Guide</i> .
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ldp statistics on page 1306 • Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain on page 576 • Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs on page 607
List of Sample Output	<p>show ldp traffic-statistics on page 1346</p> <p>show ldp traffic-statistics p2mp on page 1347</p> <p>show ldp traffic-statistics p2mp (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 1347</p> <p>show ldp traffic-statistics p2mp (Multipoint LDP with Multicast-Only Fast Reroute) on page 1348</p>

Output Fields Table 64 on page 1346 lists the output fields for the **show ldp traffic-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 64: show ldp traffic-statistics Output Fields

Field Name	Field Description
Message type	LDP message types.
FEC	Forwarding equivalence class (FEC) for which LDP traffic statistics are collected. For P2MP LSPs, FEC appears as a combination of root address and the LSP ID (root_addr:lsp_id). For M-LDP P2MP LSPs, FEC appears as a combination of root address multicast source address, and multicast group address (root_addr:lsp_id/grp,src).
Type	Type of traffic originating from a router, either Ingress (originating from this router) or Transit (forwarded through this router).
Packets	Number of packets passed by the FEC since its LSP came up.
Bytes	Number of bytes of data passed by the FEC since its LSP came up.
Shared	Whether a label is shared by prefixes: Yes or No . A Yes value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such.
Nexthop	The next hop address for P2MP LSPs. (This is the downstream LDP Session ID.)
Label	For multipoint LDP with multicast-only fast reroute (MoFRR), the multipoint LDP node selects two separate upstream peers and sends two separate labels, one to each upstream peer. The same algorithm described in RFC 6388 is used to select the primary upstream path. The backup upstream path selection again uses the same algorithm but excludes the primary upstream LSR as a candidate. Two streams of MPLS traffic are sent to the egress node from the two different upstream peers. The MPLS traffic from only one of the upstream neighbors is selected as the primary path to accept the traffic, and the other becomes the backup path. The traffic on the backup path is dropped. When the primary upstream path fails, the traffic from the backup path is then accepted. The multipoint LDP node selects the two upstream paths based on the interior gateway protocol (IGP) root node next hop. Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
Backup route	For multipoint LDP with MoFRR, the route that is used if the primary route becomes unavailable.

Sample Output

show ldp traffic-statistics

```
user@host> show ldp traffic-statistics
```

FEC	Type	Packets	Bytes	Shared
10.35.3.0/30	Transit	0	0	Yes

	Ingress	0	0	No
10.35.10.1/32	Transit	0	0	Yes
	Ingress	0	0	No
10.255.245.214/32	Transit	0	0	No
	Ingress	11	752	No
192.168.37.36/30	Transit	0	0	Yes
	Ingress	0	0	No

FEC(root_addr:lsp_id)	Nexthop	Packets	Bytes	Shared
10.255.72.160:16777217	192.168.8.81	152056	14597376	No
	192.168.8.1	152056	14597376	No
	192.168.8.65	152056	14597376	No
NET FEC Statistics:				
FEC	Type	Packets	Bytes	Shared
10.255.107.230/32	Transit	30858	2022345	No
	Ingress	20	5120	No

show ldp traffic-statistics p2mp

```
user@host> show ldp traffic-statistics p2mp
```

FEC(root_addr:lsp_id)	Nexthop	Packets	Bytes	Shared
10.255.72.160:16777217	192.168.8.81	152056	14597376	No
	192.168.8.1	152056	14597376	No
	192.168.8.65	152056	14597376	No

show ldp traffic-statistics p2mp (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show ldp traffic-statistics p2mp
```

P2MP FEC Statistics:

FEC(root_addr:lsp_id/grp,src)	Nexthop	Packets	Bytes
Shared			
11.99.0.73:239.10.0.1,11.98.0.10	11.99.0.117	243408	121217184
No			
	11.99.0.13	236286	117670428
No			
11.99.0.73:239.10.0.2,11.98.0.10	11.99.0.117	248800	123902400
No			
	11.99.0.13	240759	119897982
No			
11.99.0.73:239.10.0.1,11.98.0.20	11.99.0.117	250286	124642428
No			
	11.99.0.13	243741	121383018
No			
11.99.0.73:239.10.0.2,11.98.0.20	11.99.0.117	252970	125979060
No			
	11.99.0.13	245218	122118564
No			

show ldp traffic-statistics p2mp (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show ldp traffic-statistics p2mp
```

P2MP FEC Statistics:

FEC(root_addr:lsp_id/grp,src)	Nexthop	Packets	Bytes
Shared			
1.1.1.1:232.1.1.1,192.168.219.11, Label: 301568	1.3.8.2	0	0
No	1.3.4.2	0	0
No			
1.1.1.1:232.1.1.1,192.168.219.11, Label: 301584, Backup route	1.3.4.2	0	0
No	1.3.8.2	0	0
No			
1.1.1.1:232.1.1.2,192.168.219.11, Label: 301600	1.3.8.2	0	0
No	1.3.4.2	0	0
No			
1.1.1.1:232.1.1.2,192.168.219.11, Label: 301616, Backup route	1.3.4.2	0	0
No	1.3.8.2	0	0
No			

show security keychain

Syntax	show security keychain <brief detail>
Release Information	Command introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	Display information about authentication keychains configured for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol.
Options	none —Display information about authentication keychains. brief detail —(Optional) Display the specified level of output.
Required Privilege Level	view
List of Sample Output	show security keychain brief on page 1351 show security keychain detail on page 1351
Output Fields	Table 65 on page 1349 describes the output fields for the show security keychain command. Output fields are listed in the approximate order in which they appear.

Table 65: show security keychain Output Fields

Field Name	Field Description	Level of Output
keychain	The name of the keychain in operation.	All levels
Active-ID Send	Number of routing protocols packets sent with the active key.	All levels
Active-ID Receive	Number of routing protocols packets received with the active key.	All levels
Next-ID Send	Number of routing protocols packets sent with the next key.	All levels
Next-ID Receive	Number of routing protocols packets received with the next key.	All levels
Transition	Amount of time until the current key will be replaced with the next key in the keychain.	All levels
Tolerance	Configured clock-skew tolerance, in seconds, for accepting keys for a key chain.	All levels

Table 65: show security keychain Output Fields (*continued*)

Field Name	Field Description	Level of Output
Id	Identification number configured for the current key.	detail
Algorithm	Authentication algorithm configured for the current key.	detail
State	<p>State of the current key.</p> <p>The value can be:</p> <ul style="list-style-type: none"> • receive • send • send-receive <p>For the active key, the State can be send-receive, send, or receive. For keys that have a future start time, the State is inactive. Compare the State field to the Mode field.</p>	detail
Option	<p>For IS-IS only, the option determines how Junos OS encodes the message authentication code in routing protocol packets.</p> <p>The values can be:</p> <ul style="list-style-type: none"> • basic—Based on RFC 5304. • isis-enhanced—Based on RFC 5310. <p>The default value is basic. When you configure the isis-enhanced option, Junos OS sends RFC 5310-encoded routing protocol packets and accepts both RFC 5304-encoded and RFC 5310-encoded routing protocol packets that are received from other devices.</p> <p>When you configure basic (or do not include the options statement in the key configuration) Junos OS sends and receives RFC 5304-encoded routing protocols packets, and drops 5310-encoded routing protocol packets that are received from other devices.</p> <p>Because this setting is for IS-IS only, the TCP and the BFD protocol ignore the encoding option configured in the key.</p>	detail
Start-time	Time that the current key became active.	detail

Table 65: show security keychain Output Fields (*continued*)

Field Name	Field Description	Level of Output
Mode	<p>Mode of each key (Informational only.)</p> <p>The value can be</p> <ul style="list-style-type: none"> • receive • send • send-receive <p>The mode of the key is based on the configuration. Suppose you configure two keys, one with a start-time of today and the other with a start-time of next week. For both keys, the Mode can be send-receive, send, or receive, regardless of the configured start-time. Compare the Mode field to the State field.</p>	detail

Sample Output

show security keychain brief

```

user@host> show security keychain brief
keychain              Active-ID      Next-ID      Transition  Tolerance
                     Send  Receive    Send  Receive
hakr                   3    3           1    1         1d 23:58    3600

```

show security keychain detail

```

user@host> show security keychain detail
keychain              Active-ID      Next-ID      Transition  Tolerance
                     Send  Receive    Send  Receive
hakr                   3    3           1    1         1d 23:58    3600
  Id 3, Algorithm hmac-md5, State send-receive, Option basic
  Start-time Wed Aug 11 16:28:00 2010, Mode send-receive
  Id 1, Algorithm hmac-md5, State inactive, Option basic
  Start-time Fri Aug 20 11:30:57 2010, Mode send-receive

```

traceroute mpls ldp

Syntax `traceroute mpls <ldp> fec`
`<destination>`
`<detail>`
`<exp>`
`<fanout>`
`<logical-system>`
`<no-resolve>`
`<paths>`
`<retries>`
`<routing-instance>`
`<source>`
`<ttl>`
`<update>`
`<wait>`

Release Information Command introduced in Junos OS Release 8.4.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Trace route to a remote host for an MPLS label-switched path signaled by the LDP. Use **traceroute mpls ldp** as a debugging tool to locate MPLS label-switched path forwarding issues in a network. (Currently supported for IPv4 packets only.)

Options *fec*—Specify the IP address and optional prefix of the forwarding equivalence class (FEC).
destination—(Optional) Specify the destination address to use when sending probes.
detail—(Optional) Display detailed output.
exp—(Optional) Specify the class-of-service to use when sending probes. The range of values is 0 through 7. The default value is 7.
fanout—(Optional) Specify the maximum number of nexthops to search per node. The range of values is 1 through 16. The default value is 16.
logical-system—(Optional) Specify the name of the logical system for the traceroute attempt.
no-resolve—(Optional) Specify not to resolve the hostname that corresponds to the IP address.
paths—(Optional) Specify the number of paths to search. The range of values is 1 through 255. The default value is 16.
retries—(Optional) Specify the number of times to resend probe. values. The range of values is 1 through 9. The default value is 3.
routing-instance routing-instance-name—(Optional) Specify the name of the routing instance for the traceroute attempt.
source source-address—(Optional) Specify the source address of the outgoing traceroute packets.

ttl value—(Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds. The range of values is **1** through **125** and the default value is **64**.

wait seconds—(Optional) Specify the number of seconds to wait before resending a probe. The range of values is **5** through **15** and the default value is **10** seconds.

Required Privilege Level network

List of Sample Output [traceroute mpls ldp on page 1354](#)
[traceroute mpls ldp detail on page 1354](#)

Output Fields [Table 46 on page 1256](#) describes the output fields for the **traceroute mpls ldp fec** command and the **traceroute mpls ldp fec detail** commands. Output fields are listed in the approximate order in which they appear.

Table 66: traceroute mpls ldp Output Fields

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the traceroute mpls ldp fec command.	all levels
ttl	Time to live value of the labeled packet.	none specified
Label	Outgoing label used for forwarding the packet along the label-switched paths.	none specified
Protocol	Signaling protocol used. For this command, it is LDP.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is null .	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	detail
Parent	Address of the previous hop. Parent value for the first hop is null .	detail
Return Code	Return code for reporting the result of processing the echo request by the receiver.	detail
Response time	Time for the echo request to reach the receiver.	detail
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is none .	detail

Table 66: traceroute mpls ldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Label Stack	Label stack used to forward the packet.	detail

Sample Output

traceroute mpls ldp

```
user@router> traceroute mpls ldp 4.4.4.4
```

```
Probe options: ttl 64, retries 3, wait 10, paths 16, exp 7, fanout 16
ttl  Label Protocol Address Previous Hop Probe Status
1    100016 LDP      24.24.24.1 (null) Success
2    100000 LDP      20.20.20.2 24.24.24.1 Success
3          3 LDP      22.22.22.4 20.20.20.2 Egress
```

```
Path 1 via fe-0/3/3.101 destination 127.0.0.64
```

traceroute mpls ldp detail

```
user@router> traceroute mpls ldp 4.4.4.4 detail
```

```
Probe Options: ttl 64, retries 3, wait 10, paths 3, exp 7
Hop 24.24.24.1 Depth 1
  Parent (null)
  Return code: Label switched at stack-depth 1
  Response time 165.93 msec
  Multipath type: IP bitmask
  Address Range 1: 127.0.0.0 ~ 127.0.3.255
  Label Stack:
    Label 1 Value 100032 Protocol LDP

Hop 20.20.20.2 Depth 2
  Parent 24.24.24.1
  Return code: Upstream interface index unknown label-switched at stack-depth
1
  Response time 19.05 msec
  Multipath type: IP bitmask
  Address Range 1: 127.0.0.0 ~ 127.0.3.255
  Label Stack:
    Label 1 Value 100000 Protocol LDP

Hop 22.22.22.4 Depth 3
  Parent 20.20.20.2
  Return code: Egress-ok at stack-depth 1
  Response time 0.79 msec
  Multipath type: None
  Label Stack:
    Label 1 Value 3 Protocol LDP
```


CHAPTER 32

CCC and TCC Operational Commands

- `show connections`
- `show route ccc`
- `show route forwarding-table`

show connections

List of Syntax [Syntax on page 1356](#)
 [Syntax \(EX Series Switches\) on page 1356](#)

Syntax `show connections`
 `<brief | extensive>`
 `<all | interface-switch | lsp-switch | p2mp-receive-switch | p2mp-transmit-switch |`
 `remote-interface-switch>`
 `<down | up | up-down>`
 `<history>`
 `<labels>`
 `<logical-system (all | logical-system-name)>`
 `<name>`
 `<status>`

Syntax (EX Series Switches) `show connections`
 `<brief | extensive>`
 `<all | interface-switch | lsp-switch | p2mp-receive-switch | p2mp-transmit-switch |`
 `remote-interface-switch>`
 `<down | up | up-down>`
 `<history>`
 `<labels>`
 `<name>`
 `<status>`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.5 for EX Series switches.

Description Display information about the configured circuit cross-connect (CCC) connections.

Options **none**—Display the standard level of output for all configured CCC connections.

all—(Optional) Display all connections.

brief | extensive—(Optional) Display the specified level of output. Use history to display information about connection history. Use labels to display labels used for transmit and receive LSPs. Use status to display information about the connection and interface status.

interface-switch—(Optional) Display interface switch connections only.

lsp-switch—(Optional) Display LSP switch connections only.

p2mp-receive-switch—(Optional) Display point-to-multipoint LSP to local interfaces switch connections only.

p2mp-transmit-switch—(Optional) Display local interface to point-to-multipoint LSP switch connections only.

remote-interface-switch—(Optional) Display remote interface switch connections only.

down | up | up-down—(Optional) Display nonoperational, operational, or both kinds of connections.

history—(Optional) Display information about connection history.

labels—(Optional) Display labels used for transmit and receive.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display information about the specified connection only.

status—(Optional) Display information about the connection and interface status.

Required Privilege Level view

Output Fields [Table 22 on page 1171](#) describes the output fields for the **show connections** command. Output fields are listed in the approximate order in which they appear.

Table 67: show connections Output Fields

Field Name	Field Description
CCC and TCC connections [Link Monitoring On Off]	Whether link monitoring is enabled: On or Off .
Legend for Status (St)	Connection or circuit status. See the output's legend for an explanation of the status field values.
Legend for connection types	Type of connection: <ul style="list-style-type: none"> • if-sw—Layer 2 switching cross-connect. • rmt-if—Remote interface switch. While graceful restart is in progress, rmt-if will display a state (St) of Restart. • lsp-sw—LSP stitching cross-connect. While graceful restart is in progress, lsp-sw will display a state (St) of Restart.
Legend for circuit types	Type of circuits: <ul style="list-style-type: none"> • intf—Interface circuit. • tlsp—Transmit LSP circuit. • rlsp—Receive LSP circuit.
Connection/Circuit	Name of the configured CCC connection.
Type	Type of connection.
St	State of the connection.
Time last up	Time that the connection or circuit last transitioned to the Up (operational) state.

Table 67: show connections Output Fields (*continued*)

Field Name	Field Description
# Up trans	Number of times that the connection or circuit has transitioned to the Up (operational) state.

Sample Output

show connections

```

user@switch> show connections
CCC and TCC connections [Link Monitoring On]
  Legend for status (St)           Legend for connection types
  UN -- uninitialized             if-sw: interface switching
  NP -- not present               rmt-if: remote interface switching
  WE -- wrong encapsulation       lsp-sw: LSP switching
  DS -- disabled
  Dn -- down
  -> -- only outbound conn is up  Legend for circuit types
  <- -- only inbound conn is up   intf -- interface
  Up -- operational              tlsp -- transmit LSP
  RmtDn -- remote CCC down       rlsp -- receive LSP
  Restart -- restarting

CCC Graceful restart : Restarting

Connection/Circuit      Type   St    Time last up    # Up trans
IFSW-ed                if-sw  Up     Aug  5 15:39:15      1
  so-1/0/2.0            intf   Up
  t1-0/1/2.0            intf   Up
SW-db                  rmt-if Restart      0
  so-1/0/3.0            intf   Up
  pro4-ca                tlsp   Dn
  pro4-ac                rlsp   NP

```

show route ccc

Syntax	show route ccc ccc <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display circuit cross-connect (CCC) entries in the Multiprotocol Link Switching (MPLS) routing table.
Options	<p>ccc—Name of an entry with a circuit cross-connect interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show connections on page 1170
List of Sample Output	show route ccc extensive on page 1359
Output Fields	For information about output fields, see the output field tables for the <i>show route</i> command, the <i>show route detail</i> command, the <i>show route extensive</i> command, or the <i>show route terse</i> command.

Sample Output

show route ccc extensive

```

user@host> show route ccc fe-0/1/0.600 extensive
mpls.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
fe-0/1/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel fe-0/1/2.600.0      /16 -> {0.0.0.0}
*CCC      Preference: 7
           Next-hop reference count: 2
           Next hop: via so-0/0/3.0 weight 0x1, selected
           Label operation: Push 101424
           State: <Active Int>
           Local AS: 100
           Age: 28:13   Metric: 3
           Task: MPLS
           Announcement bits (1): 0-KRT
           AS path: I

```

show route forwarding-table

List of Syntax	Syntax on page 1360 Syntax (MX Series Routers) on page 1360 Syntax (TX Matrix and TX Matrix Plus Routers) on page 1360
Syntax	<pre>show route forwarding-table <detail extensive summary> <all> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <label name> <matching matching> <multicast> <table (default logical-system-name/routing-instance-name routing-instance-name)> <vlan (all vlan-name)> <vpn vpn></pre>
Syntax (MX Series Routers)	<pre>show route forwarding-table <detail extensive summary> <all> <bridge-domain (all domain-name)> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <label name> <learning-vlan-id learning-vlan-id> <matching matching> <multicast> <table (default logical-system-name/routing-instance-name routing-instance-name)> <vlan (all vlan-name)> <vpn vpn></pre>
Syntax (TX Matrix and TX Matrix Plus Routers)	<pre>show route forwarding-table <detail extensive summary> <all> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <matching matching> <label name> <lcc number> <multicast> <table routing-instance-name> <vpn vpn></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Option bridge-domain introduced in Junos OS Release 7.5</p> <p>Option learning-vlan-id introduced in Junos OS Release 8.4</p>

Options **all** and **vlan** introduced in Junos OS Release 9.6.

Command introduced in Junos OS Release 11.3 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.



NOTE: The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

Options **none**—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

detail | extensive | summary—(Optional) Display the specified level of output.

all—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

bridge-domain (all | bridge-domain-name)—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

ccc interface-name—(Optional) Display route entries for the specified circuit cross-connect interface.

destination destination-prefix—(Optional) Destination prefix.

family family—(Optional) Display routing table entries for the specified family: **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

interface-name interface-name—(Optional) Display routing table entries for the specified interface.

label name—(Optional) Display route entries for the specified label.

lcc number—(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers, display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

learning-vlan-id *learning-vlan-id*—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

matching *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast—(Optional) Display routing table entries for multicast routes.

table (**default** | *logical-system-name/routing-instance-name* | *routing-instance-name*)—(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the **show route instance** command.

vlan (**all** | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

vpn *vpn*—(Optional) Display routing table entries for a specified VPN.

Required Privilege Level

view

List of Sample Output

[show route forwarding-table on page 1365](#)
[show route forwarding-table detail on page 1366](#)
[show route forwarding-table destination extensive \(Weights and Balances\) on page 1366](#)
[show route forwarding-table extensive on page 1367](#)
[show route forwarding-table extensive \(RPF\) on page 1368](#)
[show route forwarding-table family mpls on page 1369](#)
[show route forwarding-table family vpls on page 1369](#)
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled\) on page 1369](#)
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled with MAC Statistics\) on page 1370](#)
[show route forwarding-table family vpls extensive on page 1370](#)
[show route forwarding-table table default on page 1371](#)
[show route forwarding-table table logical-system-name/routing-instance-name on page 1372](#)

[show route forwarding-table vpn on page 1373](#)

Output Fields [Table 68 on page 1363](#) lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

Table 68: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Logical system	Name of the logical system. This field is displayed if you specify the table logical-system-name/routing-instance-name option on a device that is configured for and supports logical systems.	All levels
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels
Address family	Address family (for example, IP, IPv6, ISO, MPLS, and VPLS).	All levels
Destination	Destination of the route.	detail extensive
Route Type (Type)	How the route was placed into the forwarding table. When the detail keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> cloned (clon)—(TCP or multicast only) Cloned route. destination (dest)—Remote addresses directly reachable through an interface. destination down (iddn)—Destination route for which the interface is unreachable. interface cloned (ifcl)—Cloned route for which the interface is unreachable. route down (ifdn)—Interface route for which the interface is unreachable. ignore (ignr)—Ignore this route. interface (intf)—Installed as a result of configuring an interface. permanent (perm)—Routes installed by the kernel when the routing table is initialized. user—Routes installed by the routing protocol process or as a result of the configuration. 	All levels
Route Reference (RtRef)	Number of routes to reference.	detail extensive
Flags	Route type flags: <ul style="list-style-type: none"> none—No flags are enabled. accounting—Route has accounting enabled. cached—Cache route. incoming-iface interface-number—Check against incoming interface. prefix load balance—Load balancing is enabled for this prefix. rt nh decoupled—Route has been decoupled from the next hop to the destination. sent to PFE—Route has been sent to the Packet Forwarding Engine. static—Static route. 	extensive
Next hop	IP address of the next hop to the destination.	detail extensive

Table 68: show route forwarding-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
Next hop Type (Type)	<p>Next-hop type. When the detail keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> • broadcast (bcst)—Broadcast. • deny—Deny. • discard (dscd)—Discard. • hold—Next hop is waiting to be resolved into a unicast or multicast type. • indexed (idxd)—Indexed next hop. • indirect (indr)—Indirect next hop. • local (locl)—Local address on an interface. • routed multicast (mcrst)—Regular multicast next hop. • multicast (mcst)—Wire multicast next hop (limited to the LAN). • multicast discard (mdsc)—Multicast discard. • multicast group (mgrp)—Multicast group member. • receive (rcv)—Receive. • reject (rjct)—Discard. An ICMP unreachable message was sent. • resolve (rslv)—Resolving the next hop. • unicast (ucst)—Unicast. • unilist (ulst)—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list. 	detail extensive
Index	Software index of the next hop that is used to route the traffic for a given prefix.	detail extensive none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	extensive
Reference (NhRef)	Number of routes that refer to this next hop.	detail extensive none
Next-hop interface (Netif)	Interface used to reach the next hop.	detail extensive none
Weight	Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the Balance field description).	extensive
Balance	Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.	extensive
RPF interface	List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when rpf-check is configured on the interface.	extensive

Sample Output

show route forwarding-table

```

user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  46   4
0.0.0.0/32       perm  0                               dscd  44   1
1.1.1.0/24       ifdn  0                               rslv  608  1 ge-2/0/1.0
1.1.1.0/32       iddn  0 1.1.1.0             recv  606  1 ge-2/0/1.0
1.1.1.1/32       user  0                               rjct  46   4
1.1.1.1/32       intf  0 1.1.1.1             locl  607  2
1.1.1.1/32       iddn  0 1.1.1.1             locl  607  2
1.1.1.255/32     iddn  0 ff:ff:ff:ff:ff:ff    bcst  605  1 ge-2/0/1.0
10.0.0.0/24      intf  0                               rslv  616  1 ge-2/0/0.0
10.0.0.0/32      dest  0 10.0.0.0             recv  614  1 ge-2/0/0.0
10.0.0.1/32      intf  0 10.0.0.1             locl  615  2
10.0.0.1/32      dest  0 10.0.0.1             locl  615  2
10.0.0.255/32    dest  0 10.0.0.255          bcst  613  1 ge-2/0/0.0
10.1.1.0/24      ifdn  0                               rslv  612  1 ge-2/0/1.0
10.1.1.0/32      iddn  0 10.1.1.0             recv  610  1 ge-2/0/1.0
10.1.1.1/32      user  0                               rjct  46   4
10.1.1.1/32      intf  0 10.1.1.1             locl  611  2
10.1.1.1/32      iddn  0 10.1.1.1             locl  611  2
10.1.1.255/32    iddn  0 ff:ff:ff:ff:ff:ff    bcst  609  1 ge-2/0/1.0
10.209.0.0/16    user  0 10.209.63.254        ucst  419  20 fxp0.0
10.209.0.0/16    user  1 0:12:1e:ca:98:0      ucst  419  20 fxp0.0
10.209.0.0/18    intf  0                               rslv  418  1 fxp0.0
10.209.0.0/32    dest  0 10.209.0.0           recv  416  1 fxp0.0
10.209.2.131/32  intf  0 10.209.2.131         locl  417  2
10.209.2.131/32  dest  0 10.209.2.131         locl  417  2
10.209.17.55/32  dest  0 0:30:48:5b:78:d2     ucst  435  1 fxp0.0
10.209.63.42/32  dest  0 0:23:7d:58:92:ca     ucst  434  1 fxp0.0
10.209.63.254/32 dest  0 0:12:1e:ca:98:0      ucst  419  20 fxp0.0
10.209.63.255/32 dest  0 10.209.63.255        bcst  415  1 fxp0.0
10.227.0.0/16    user  0 10.209.63.254        ucst  419  20 fxp0.0

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  27   1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf  0                               locl  28   1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  6    1
ff00::/8         perm  0                               mdsc  4    1
ff02::1/128      perm  0 ff02::1             mcst  3    1

Routing table: ccc
MPLS:
Interface.Label  Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  16   1
100004(top)fe-0/0/1.0

```

show route forwarding-table detail

```

user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          user   2 0:90:69:8e:b1:1b ucst  132   4 fxp0.0
default          perm   0                               rjct   14    1
10.1.1.0/24      intf   0 ff.3.0.21         ucst  322   1 so-5/3/0.0
10.1.1.0/32      dest   0 10.1.1.0          recv  324   1 so-5/3/0.0
10.1.1.1/32      intf   0 10.1.1.1          locl  321   1
10.1.1.255/32    dest   0 10.1.1.255        bcst  323   1 so-5/3/0.0
10.21.21.0/24    intf   0 ff.3.0.21         ucst  326   1 so-5/3/0.0
10.21.21.0/32    dest   0 10.21.21.0        recv  328   1 so-5/3/0.0
10.21.21.1/32    intf   0 10.21.21.1        locl  325   1
10.21.21.255/32  dest   0 10.21.21.255      bcst  327   1 so-5/3/0.0
127.0.0.1/32     intf   0 127.0.0.1         locl  320   1
172.17.28.19/32  clon   1 192.168.4.254     ucst  132   4 fxp0.0
172.17.28.44/32  clon   1 192.168.4.254     ucst  132   4 fxp0.0

...

Routing table: private1__inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct   46    1
10.0.0.0/8       intf   0                               rslv  136   1 fxp1.0
10.0.0.0/32      dest   0 10.0.0.0          recv  134   1 fxp1.0
10.0.0.4/32      intf   0 10.0.0.4          locl  135   2
10.0.0.4/32      dest   0 10.0.0.4          locl  135   2

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct   38    1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct   22    1
ff00::/8         perm   0                               mdsc   21    1
ff02::1/128      perm   0 ff02::1          mcst   17    1

...

Routing table: mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct  28    1

```

show route forwarding-table destination extensive (Weights and Balances)

```

user@host> show route forwarding-table destination 3.4.2.1 extensive
Routing table: inet [Index 0]
Internet:

Destination: 3.4.2.1/32
Route type: user
Route reference: 0                               Route interface-index: 0

```

```

Flags: sent to PFE
Next-hop type: unilist           Index: 262143  Reference: 1
Nexthop: 4.4.4.4
Next-hop type: unicast          Index: 335      Reference: 2
Next-hop interface: so-1/1/0.0  Weight: 22    Balance: 3
Nexthop: 145.12.1.2
Next-hop type: unicast          Index: 337      Reference: 2
Next-hop interface: so-0/1/2.0  Weight: 33    Balance: 33

```

show route forwarding-table extensive

```

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:

Destination: default
Route type: user
Route reference: 2                Route interface-index: 0
Flags: sent to PFE
Nexthop: 0:90:69:8e:b1:1b
Next-hop type: unicast           Index: 132      Reference: 4
Next-hop interface: fxp0.0

Destination: default
Route type: permanent
Route reference: 0                Route interface-index: 0
Flags: none
Next-hop type: reject            Index: 14       Reference: 1

Destination: 127.0.0.1/32
Route type: interface
Route reference: 0                Route interface-index: 0
Flags: sent to PFE
Nexthop: 127.0.0.1
Next-hop type: local              Index: 320      Reference: 1

...

Routing table: private1__inet [Index 1]
Internet:

Destination: default
Route type: permanent
Route reference: 0                Route interface-index: 0
Flags: sent to PFE
Next-hop type: reject            Index: 46       Reference: 1

Destination: 10.0.0.0/8
Route type: interface
Route reference: 0                Route interface-index: 3
Flags: sent to PFE
Next-hop type: resolve           Index: 136      Reference: 1
Next-hop interface: fxp1.0

...

Routing table: iso [Index 0]
ISO:

Destination: default
Route type: permanent

```

```

Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 38      Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 22      Reference: 1

Destination: ff00::/8
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: multicast discard
Route interface-index: 0
Index: 21      Reference: 1

...

Routing table: private1__inet6 [Index 1]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 54      Reference: 1

Destination: fe80::2a0:a5ff:fe3d:375/128
Route type: interface
Route reference: 0
Flags: sent to PFE
Next-hop: fe80::2a0:a5ff:fe3d:375
Next-hop type: local
Route interface-index: 0
Index: 75      Reference: 1

...

```

show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```

so-1/1/0 {
  unit 0 {
    family inet {
      rpf-check;
      address 15.95.1.2/30;
    }
  }
}

```

```

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:
...
...
Destination: 15.95.1.3/32
Route type: destination
Route reference: 0
Route interface-index: 67

```

```

Flags: sent to PFE
Nexthop: 15.95.1.3
Next-hop type: broadcast           Index: 328       Reference: 1
Next-hop interface: so-1/1/0.0
RPF interface: so-1/1/0.0

```

show route forwarding-table family mpls

```

user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
0                user  0
1                user  0
2                user  0
100000           user  0 10.31.1.6          swap 100001 fe-1/1/0.0
800002           user  0                  Pop          vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
                  user  0                  indr  351    4
                  Push 800000, Push 100002(top)

so-0/0/0.0

```

show route forwarding-table family vpls

```

user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          dymn  0
default          perm  0
fe-0/1/0.0       dymn  0
00:90:69:0c:20:1f/48 <<<<<Remote CE
                  dymn  0                  indr  351    4
                  Push 800000, Push 100002(top)

so-0/0/0.0
00:90:69:85:b0:1f/48 <<<<<Local CE
                  dymn  0                  ucst  354    2 fe-0/1/0.0

```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled)

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
lsi.1048832      intf  0
                  4.4.3.2          indr 1048574 4
                  Push 262145      621    2

ge-3/0/0.0
00:19:e2:25:d0:01/48 user  0                  ucst  590    5 ge-2/3/9.0
0x30003/51       user  0                  comp  627    2
ge-2/3/9.0       intf  0                  ucst  590    5 ge-2/3/9.0
ge-3/1/3.0       intf  0                  ucst  619    4 ge-3/1/3.0
0x30002/51       user  0                  comp  600    2
0x30001/51       user  0                  comp  597    2

```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics)

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats

```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		dscd	519	1	
1si.1048834	intf	0		indr	1048574	4	
			4.4.3.2	Push	262145	592	2
ge-3/0/0.0							
00:19:e2:25:d0:01/48	user	0		ucst	590	5	ge-2/3/9.0
0x30003/51	user	0		comp	630	2	
ge-2/3/9.0	intf	0		ucst	590	5	ge-2/3/9.0
ge-3/1/3.0	intf	0		ucst	591	4	ge-3/1/3.0
0x30002/51	user	0		comp	627	2	
0x30001/51	user	0		comp	624	2	

show route forwarding-table family vpls extensive

```

user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

```

Destination: default

Route type: dynamic	Route interface-index: 72
Route reference: 0	
Flags: sent to PFE	
Next-hop type: flood	Index: 289 Reference: 1
Next-hop type: unicast	Index: 291 Reference: 3
Next-hop interface: fe-0/1/3.0	
Next-hop type: unicast	Index: 290 Reference: 3
Next-hop interface: fe-0/1/2.0	

Destination: default

Route type: permanent	Route interface-index: 0
Route reference: 0	
Flags: none	
Next-hop type: discard	Index: 341 Reference: 1

Destination: fe-0/1/2.0

Route type: dynamic	Route interface-index: 69
Route reference: 0	
Flags: sent to PFE	
Next-hop type: flood	Index: 293 Reference: 1
Next-hop type: indirect	Index: 363 Reference: 4
Next-hop type: Push 800016	
Next-hop interface: at-1/0/1.0	
Next-hop type: indirect	Index: 301 Reference: 5
Next hop: 10.31.3.2	
Next-hop type: Push 800000	
Next-hop interface: fe-0/1/1.0	
Next-hop type: unicast	Index: 291 Reference: 3
Next-hop interface: fe-0/1/3.0	

Destination: fe-0/1/3.0

Route type: dynamic	Route interface-index: 70
Route reference: 0	
Flags: sent to PFE	
Next-hop type: flood	Index: 292 Reference: 1


```

Next-hop type: indirect          Index: 363      Reference: 4
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect          Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast           Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0

Destination: 10:00:00:01:01:01/48
Route type: dynamic
Route reference: 0               Route interface-index: 70
Flags: sent to PFE, prefix load balance
Next-hop type: unicast           Index: 291      Reference: 3
Next-hop interface: fe-0/1/3.0
Route used as destination:
  Packet count:      6640    Byte count:      675786
Route used as source
  Packet count:      6894    Byte count:      696424

Destination: 10:00:00:01:01:04/48
Route type: dynamic
Route reference: 0               Route interface-index: 69
Flags: sent to PFE, prefix load balance
Next-hop type: unicast           Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0
Route used as destination:
  Packet count:      96     Byte count:      8079
Route used as source:
  Packet count:      296    Byte count:      24955

Destination: 10:00:00:01:03:05/48
Route type: dynamic
Route reference: 0               Route interface-index: 74
Flags: sent to PFE, prefix load balance
Next-hop type: indirect          Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0

```

show route forwarding-table table default

```

user@host> show route forwarding-table table default
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
0.0.0.0/32       perm  0
10.0.60.0/30     user  0 10.0.60.13          ucst  713  5 fe-0/1/3.0
10.0.60.12/30    intf  0                   rslv  688  1 fe-0/1/3.0
10.0.60.12/32    dest  0 10.0.60.12          recv  686  1 fe-0/1/3.0
10.0.60.13/32    dest  0 0:5:85:8b:bc:22     ucst  713  5 fe-0/1/3.0
10.0.60.14/32    intf  0 10.0.60.14          locl  687  2
10.0.60.14/32    dest  0 10.0.60.14          locl  687  2
10.0.60.15/32    dest  0 10.0.60.15          bcst  685  1 fe-0/1/3.0
10.0.67.12/30    user  0 10.0.60.13          ucst  713  5 fe-0/1/3.0
10.0.80.0/30     ifdn  0 ff.3.0.21          ucst  676  1 so-0/0/1.0
10.0.80.0/32     dest  0 10.0.80.0           recv  678  1 so-0/0/1.0
10.0.80.2/32     user  0                   rjct  36   2
10.0.80.2/32     intf  0 10.0.80.2           locl  675  1

```

```

10.0.80.3/32      dest    0 10.0.80.3      bcst   677    1 so-0/0/1.0
10.0.90.12/30     intf    0                rslv   684    1 fe-0/1/0.0
10.0.90.12/32     dest    0 10.0.90.12    recv   682    1 fe-0/1/0.0
10.0.90.14/32     intf    0 10.0.90.14     locl   683    2
10.0.90.14/32     dest    0 10.0.90.14     locl   683    2
10.0.90.15/32     dest    0 10.0.90.15     bcst   681    1 fe-0/1/0.0
10.5.0.0/16       user    0 192.168.187.126 ucst   324   15 fxp0.0
10.10.0.0/16      user    0 192.168.187.126 ucst   324   15 fxp0.0
10.13.10.0/23     user    0 192.168.187.126 ucst   324   15 fxp0.0
10.84.0.0/16      user    0 192.168.187.126 ucst   324   15 fxp0.0
10.150.0.0/16     user    0 192.168.187.126 ucst   324   15 fxp0.0
10.157.64.0/19    user    0 192.168.187.126 ucst   324   15 fxp0.0
10.209.0.0/16     user    0 192.168.187.126 ucst   324   15 fxp0.0

```

...

Routing table: default.iso

ISO:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	60	1	

Routing table: default.inet6

Internet6:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	44	1	
::/128	perm	0		dscd	42	1	
ff00::/8	perm	0		mdsc	43	1	
ff02::1/128	perm	0	ff02::1	mcst	39	1	

Routing table: default.mpls

MPLS:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		dscd	50	1	

show route forwarding-table table logical-system-name/routing-instance-name

```
user@host> show route forwarding-table table R4/vpn-red
```

Logical system: R4

Routing table: vpn-red.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	563	1	
0.0.0.0/32	perm	0		dscd	561	2	
1.0.0.1/32	user	0		dscd	561	2	
2.0.2.0/24	intf	0		rslv	771	1	ge-1/2/0.3
2.0.2.0/32	dest	0	2.0.2.0	recv	769	1	ge-1/2/0.3
2.0.2.1/32	intf	0	2.0.2.1	locl	770	2	
2.0.2.1/32	dest	0	2.0.2.1	locl	770	2	
2.0.2.2/32	dest	0	0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0	ucst	789	1	ge-1/2/0.3
2.0.2.255/32	dest	0	2.0.2.255	bcst	768	1	ge-1/2/0.3
224.0.0.0/4	perm	1		mdsc	562	1	
224.0.0.1/32	perm	0	224.0.0.1	mcst	558	1	
255.255.255.255/32	perm	0		bcst	559	1	

Logical system: R4

Routing table: vpn-red.iso

ISO:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	608	1	

```

Logical system: R4
Routing table: vpn-red.inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct  708   1
::/128           perm  0                dscd  706   1
ff00::/8         perm  0                mdsc  707   1
ff02::1/128     perm  0 ff02::1          mcst  704   1

```

```

Logical system: R4
Routing table: vpn-red.mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                dscd  638

```

show route forwarding-table vpn

```

user@host> show route forwarding-table vpn VPN-A
Routing table:: VPN-A.inet
Internet:
Destination      Type RtRef Nexthop          Type Index NhRef Netif
default          perm  0                rjct   4    4
10.39.10.20/30   intf  0 ff.3.0.21          ucst   40    1
so-0/0/0.0
10.39.10.21/32   intf  0 10.39.10.21        locl   36    1
10.255.14.172/32 user   0                ucst   69    2
so-0/0/0.0
10.255.14.175/32 user   0                indr   81    3
Push 100004, Push
100004(top) so-1/0/0.0
224.0.0.0/4      perm  2                mdsc   5    3
224.0.0.1/32     perm  0 224.0.0.1          mcst   1    8
224.0.0.5/32     user  1 224.0.0.5          mcst   1    8
255.255.255.255/32 perm  0                bcst   2    3

```


CHAPTER 33

PCEP Operational Commands

- `clear path-computation-client statistics`
- `request path-computation-client active-pce`
- `show path-computation-client active-pce`
- `show path-computation-client statistics`

clear path-computation-client statistics

Syntax	clear path-computation-client statistics < <i>pce-id</i> all>
Release Information	Command introduced in Junos OS Release 12.3.
Description	Clear Path Computation Element (PCE) statistics.
Options	<i>pce-id</i> —(Optional) Clear statistics of the specified PCE. all—(Optional) Clear statistics of all available PCEs configured on the path computation client (PCC).
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show path-computation-client statistics on page 1382
List of Sample Output	clear path-computation-client statistics pce-id on page 1376 clear path-computation-client statistics all on page 1376
Output Fields	When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear path-computation-client statistics pce-id

```
user@host> clear path-computation-client statistics pce1
```

clear path-computation-client statistics all

```
user@host> clear path-computation-client statistics all
```

request path-computation-client active-pce

Syntax	<code>request path-computation-client active-pce <i>pce-id</i></code>
Release Information	Command introduced in Junos OS Release 12.3.
Description	Request a new active Path Computation Element (PCE).
Options	<i>pce-id</i> —Unique user defined ID for this PCE.
Required Privilege Level	request
Related Documentation	<ul style="list-style-type: none">• show path-computation-client active-pce on page 1378
List of Sample Output	request path-computation-client active-pce pce-id on page 1377

Sample Output

`request path-computation-client active-pce pce-id`

```
user@host> request path-computation-client active-pce pce1
```

show path-computation-client active-pce

Syntax	show path-computation-client active-pce <brief detail>
Release Information	Command introduced in Junos OS Release 12.3.
Description	Displays information about the current active Path Computation Element (PCE).
Options	none —Display brief information about the current active PCE. brief detail —(Optional) Display the specific level of output.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • request path-computation-client active-pce on page 1377
List of Sample Output	show path-computation-client active-pce on page 1380 show path-computation-client active-pce detail on page 1380
Output Fields	Table 69 on page 1378 describes the output fields for the show path-computation-client active-pce command. Output fields are listed in the approximate order in which they appear.

Table 69: show path-computation-client active-pce Output Fields

Field Name	Field Description	Level of Output
IP address	IP address of the current active PCE.	All levels
Priority	Active PCE priority.	All levels
PCE status	Active PCE state: <ul style="list-style-type: none"> • PCE_STATE_NEW—Initial PCEP session state. • PCE_STATE_RECONNECT—Trying to re-establish TCP connection with the PCEP peer. • PCE_STATE_CONNECTING—Establishing TCP connection with the PCEP peer. • PCE_STATE_CONNECTED—TCP connection established with the PCEP peer. • PCE_STATE_SYNC—Open messages exchanged with the PCEP peer and entering SYNC state. • PCE_STATE_UP—PCEP session established. 	All levels
Session type	Active PCE type: <ul style="list-style-type: none"> • PCE_TYPE_STATELESS—Does not learn LSP state information from PCC. • PCE_TYPE_STATEFUL—Uses LSP state information learned from PCCs to optimize path computations, but does not actively update LSP state. A PCC maintains synchronization with the PCE. • PCE_TYPE_STATEFULACTIVE—Uses LSP state information learned from PCCs to optimize path computations, and actively updates LSP parameters in those PCCs that delegate control of their LSPs to the PCE. 	All levels

Table 69: show path-computation-client active-pce Output Fields (*continued*)

Field Name	Field Description	Level of Output
PCE-mastership	PCE mastership state: <ul style="list-style-type: none"> main—Current active PCE. backup—Backup PCE. 	All levels
PCRpts	Number of PC report (PCRpt) messages sent by PCC to a stateful PCE to report current state of LSP(s).	All levels
PCUpdates	Number of PC update (PCUpd) messages sent by a PCE to a PCC to update LSP parameters.	All levels
Local Keepalive timer	Keepalive timer used by or for the PCC.	All levels
Local Dead timer	Dead timer used by or for the PCC.	All levels
Remote Keepalive timer	Keepalive timer used by or for the PCE.	All levels
Remote Dead timer	Dead timer used by or for the PCE.	All levels
PCErr-recv	Information about type, value, and number of PC Error messages received.	All levels
Max unknown messages	Maximum number of unknown messages received for a PCEP session. Recommended value is 5. If the number of unknown messages received by a PCC or PCE is greater than or equal to the maximum number, the PCEP session is closed.	detail
Keepalives received	Number of Keepalive messages received by a PCC from a PCE.	detail
Keepalives sent	Number of Keepalive messages sent by a PCC to a PCE.	detail
Dead timer	Dead timer used by the current active PCE.	detail
Elapsed as main current	Time (in seconds) the PCE is in the main mastership state.	detail
Elapsed as main total	Time (in seconds) the PCE became main from the last PCCD restart.	detail
Unknown msgs/min rate	Number of unknown messages received per minute.	detail
Session failures	Number of PCEP session failures with the PCE.	detail
Delegation timeout in	Time (in seconds) left for LSP delegation to timeout.	detail
Delegation failures	Number of LSP delegation failures.	detail

Table 69: show path-computation-client active-pce Output Fields (*continued*)

Field Name	Field Description	Level of Output
Connection down	Time (in seconds) since the PCEP session is down.	detail
PCErr-sent	Information about type, value, and number of PC Error messages sent.	All levels

Sample Output

show path-computation-client active-pce

```

user@host> show path-computation-client active-pce
PCE pce1
General
  IP address           : 10.209.57.166
  Priority              : 2
  PCE status           : PCE_STATE_NEW
  Session type         : PCE_TYPE_STATEFULACTIVE
  PCE-mastership       : main

Counters
  PCReqs               Total: 0          last 5min: 0          last hour: 0
  PCReps               Total: 0          last 5min: 0          last hour: 0
  PCRpts              Total: 0          last 5min: 0          last hour: 0
  PCUpdates            Total: 0          last 5min: 0          last hour: 0

Timers
  Local                Keepalive timer: 0 [s]   Dead timer: 0 [s]
  Remote               Keepalive timer: 0 [s]   Dead timer: 0 [s]

Errors
  PCErr-recv
  PCErr-sent
    Type: 19          Value: 3          Count: 1
  PCE-PCC-NTFS
  PCC-PCE-NTFS

```

show path-computation-client active-pce detail

```

user@host> show path-computation-client active-pce detail
PCE pce1
General
  IP address           : 172.22.25.223
  Priority              : 1
  PCE status           : PCE_STATE_RECONNECT
  Session type         : PCE_TYPE_STATEFULACTIVE
  PCE-mastership       : main
  Max unknown messages : 5
  Keepalives received  : 0
  Keepalives sent      : 0
  Dead timer           : 0 [s]
  Elapsed as main current : 1 [s]

```

```

Elapsed as main total : 2542 [s]
Unknown msgs/min rate : 0
Session failures      : 575
Delegation timeout in : 14 [s]
Delegation failures   : 21928
Connection down       : 16 [s]

```

Counters

```

PCReqs          Total: 0          last 5min: 0          last hour: 0

PCReps          Total: 0          last 5min: 0          last hour: 0

PCRpts          Total: 31512       last 5min: 7243       last hour:
7243
PCUpdates       Total: 80          last 5min: 40          last hour:
40

```

Timers

```

Local           Keepalive timer: 30 [s]  Dead timer: 120 [s]

Remote          Keepalive timer: 30 [s]  Dead timer: 120 [s]

```

Errors

```

PCErr-recv
PCErr-sent
Type: 1          Value: 2          Count: 12

PCE-PCC-NTFS
PCC-PCE-NTFS

```

show path-computation-client statistics

Syntax	show path-computation-client statistics <brief detail> <all>
Release Information	Command introduced in Junos OS Release 12.3.
Description	Display statistics about the Path Computation Element (PCE).
Options	<p>none—Display statistics about the primary PCE.</p> <p>brief detail—(Optional) Display the specific level of output.</p> <p>all—(Optional) Display the statistics about all PCEs configured on the PCC.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear path-computation-client statistics on page 1376
List of Sample Output	show path-computation-client statistics all on page 1384 show path-computation-client statistics detail on page 1385
Output Fields	Table 70 on page 1382 describes the output fields for the show path-computation-client statistics command. Output fields are listed in the approximate order in which they appear.

Table 70: show path-computation-client statistics Output Fields

Field Name	Field Description	Level of Output
IP address	IP address of the PCE.	All levels
Priority	PCE priority.	All levels
PCE status	PCE state: <ul style="list-style-type: none"> PCE_STATE_NEW— Initial PCEP session state. PCE_STATE_RECONNECT—Trying to re-establish TCP connection with the PCEP peer. PCE_STATE_CONNECTING—Establishing TCP connection with the PCEP peer. PCE_STATE_CONNECTED—TCP connection established with the PCEP peer. PCE_STATE_SYNC—Open messages exchanged with the PCEP peer and entering SYNC state. PCE_STATE_UP—PCEP session established. 	All levels

Table 70: show path-computation-client statistics Output Fields (continued)

Field Name	Field Description	Level of Output
Session type	Active PCE type: <ul style="list-style-type: none"> • PCE_TYPE_STATELESS—Does not learn LSP state information from PCC. • PCE_TYPE_STATEFUL—Uses LSP state information learned from PCCs to optimize path computations, but does not actively update LSP state. A PCC maintains synchronization with the PCE. • PCE_TYPE_STATEFULACTIVE—Uses LSP state information learned from PCCs to optimize path computations, and actively updates LSP parameters in those PCCs that delegate control of their LSPs to the PCE. 	All levels
PCE-mastership	PCE mastership state: <ul style="list-style-type: none"> • main • primary • backup 	All levels
PCRpts	Number of PC report (PCRpt) messages sent by PCC to a stateful PCE to report current state of LSP(s).	All levels
PCUpdates	Number of PC update (PCUpd) messages sent by a PCE to a PCC to update LSP parameters.	All levels
Local Keepalive timer	Keepalive timer used by or for the PCC.	All levels
Local Dead timer	Dead timer used by or for the PCC.	All levels
Remote Keepalive timer	Keepalive timer used by or for the PCE.	All levels
Remote Dead timer	Dead timer used by or for the PCE.	All levels
PCErr-recv	Information about type, value, and number of PC Error messages received.	All levels
PCErr-sent	Information about type, value, and number of PC Error messages sent.	All levels
Max unknown messages	Maximum number of unknown messages received for a PCEP session. Recommended value is 5. If the number of unknown messages received by a PCC or PCE is greater than or equal to the maximum number, the PCEP session is closed.	detail

Table 70: show path-computation-client statistics Output Fields (continued)

Field Name	Field Description	Level of Output
Keepalives received	Number of Keepalive messages received by a PCC from a PCE.	detail
Keepalives sent	Number of Keepalive messages sent by a PCC to a PCE.	detail
Elapsed as main current	Time (in seconds) the PCE is in the main mastership state.	detail
Elapsed as main total	Time (in seconds) the PCE became main from the last PCCD restart.	detail
Unknown msgs/min rate	Number of unknown messages received per minute.	detail
Session failures	Number of PCEP session failures with the PCE.	detail
Delegation timeout in	Time (in seconds) left for LSP delegation to timeout.	detail
Delegation failures	Number of LSP delegation failures.	detail
Connection down	Time (in seconds) since the PCEP session is down.	detail

Sample Output

show path-computation-client statistics all

```

user@host> show path-computation-client statistics all
PCE pce1

General
  IP address       : 10.209.57.166
  Priority          : 2
  PCE status       : PCE_STATE_NEW
  Session type     : PCE_TYPE_STATEFULACTIVE
  PCE-mastership   : main

Counters
  PCReqs          Total: 0          last 5min: 0          last hour: 0
  PCReps          Total: 0          last 5min: 0          last hour: 0
  PCRpts          Total: 0          last 5min: 0          last hour: 0
  PCUpdates       Total: 0          last 5min: 0          last hour: 0

Timers
  Local          Keepalive timer:      0 [s]   Dead timer:      0 [s]

```

```

Remote           Keepalive timer:      0 [s]   Dead timer:      0 [s]

Errors
  PCErr-recv
  PCErr-sent
  PCE-PCC-NTFS
  PCC-PCE-NTFS

PCE pce2

General
  IP address      : 10.31.32.1
  Priority        : 10
  PCE status      : PCE_STATE_NEW
  Session type    : PCE_TYPE_STATEFULACTIVE
  PCE-mastership  : backup

Counters
  PCReqs          Total: 0           last 5min: 0       last hour: 0
  PCReps          Total: 0           last 5min: 0       last hour: 0
  PCRpts          Total: 0           last 5min: 0       last hour: 0
  PCUpdates       Total: 0           last 5min: 0       last hour: 0

Timers
  Local           Keepalive timer:      0 [s]   Dead timer:      0 [s]
  Remote          Keepalive timer:      0 [s]   Dead timer:      0 [s]

Errors
  PCErr-recv
  PCErr-sent
  PCE-PCC-NTFS
  PCC-PCE-NTFS

```

show path-computation-client statistics detail

```

user@host> show path-computation-client statistics detail
PCE pce1
General
  IP address      : 10.209.57.166
  Priority        : 2
  PCE status      : PCE_STATE_NEW
  Session type    : PCE_TYPE_STATEFULACTIVE
  PCE-mastership  : main
  Max unknown messages : 5
  Keepalives received : 0
  Keepalives sent    : 0
  Dead timer       : 0 [s]
  Elapsed as main current : 294 [s]
  Elapsed as main total   : 294 [s]
  Unknown msgs/min rate  : 0
  Session failures      : 0
  Replies timedout      : 0
  Delegation timeout in : 26 [s]
  Delegation failures    : 0
  Connection down       : 4 [s]

```

Counters

PCReqs	Total: 0	last 5min: 0	last hour: 0
PCReps	Total: 0	last 5min: 0	last hour: 0
PCRpts	Total: 0	last 5min: 0	last hour: 0
PCUpdates	Total: 0	last 5min: 0	last hour: 0

Timers

Local	Keepalive timer:	0 [s]	Dead timer:	0 [s]
Remote	Keepalive timer:	0 [s]	Dead timer:	0 [s]

Errors

PCErr-recv
PCErr-sent
PCE-PCC-NTFS
PCC-PCE-NTFS

PART 11

Index

- [Index on page 1389](#)

Index

Symbols

#, comments in configuration statements.....	xl
(), in syntax descriptions.....	xl
< >, in syntax descriptions.....	xl
[], in configuration statements.....	xl
{ }, in configuration statements.....	xl
(pipe), in syntax descriptions.....	xl

A

activate interface command	794
adaptive rerouting.....	249, 825
adaptive statement.....	825
usage guidelines.....	249
address (tracing flag).....	1078
address statement	
LMP	1102
usage guidelines.....	681
addresses	
associating with LSPs.....	228, 880
egress router address.....	212, 953
ingress router.....	212, 872
adjust-interval statement.....	826
usage guidelines.....	259
adjust-threshold statement.....	826
usage guidelines.....	260
adjust-threshold-overflow-limit statement.....	827
usage guidelines.....	261
adjust-threshold-underflow-limit statement.....	828
usage guidelines.....	261
Admin Status object, GMPLS.....	690
admin-down statement.....	828
configuration guidelines.....	690
admin-group statement	
bypass LSPs.....	973
configuration	
bypass LSPs.....	504
LSPs.....	829
MPLS interfaces.....	829
admin-group-extended statement.....	830
usage guidelines.....	242

admin-groups statement.....	831
usage guidelines.....	240
admin-groups-extended statement.....	832
usage guidelines.....	242
admin-groups-extended-range statement.....	833
usage guidelines.....	242
administrative groups See groups	
admin-groups statement.....	831
configuration.....	240
exclude statement.....	861
extended.....	242
fast reroute.....	226
include-all statement.....	877
include-any statement.....	878
administrative groups, MPLS	760
advertise-mode statement	
MPLS.....	834
advertisement messages, LDP.....	522
advertisement-hold-time statement.....	835
usage guidelines.....	268
aggregate statement	
RSVP.....	974
usage guidelines.....	473
aggregated Ethernet interfaces.....	650
aggregated interfaces.....	39
aggregation, RSVP.....	974
all (tracing flag).....	954
LMP	1120
RSVP.....	1012
allocation of labels.....	26
allow-fragmentation statement.....	835
usage guidelines.....	484
allow-subnet-mismatch statement.....	1023
usage guidelines.....	633
always-mark-connection-protection-tlv	
statement.....	836
usage guidelines.....	479
ARP configuration.....	662
associate-backup-pe-groups statement.....	836
usage guidelines.....	307
associate-lsp	
usage guidelines.....	136
associate-lsp statement	
MPLS-TP.....	837
Associated bi-directional LSP.....	52
ATM	
circuits.....	648, 656
ATM encapsulation	
Layer 2 TCC.....	660

authentication		
RSVP.....	476, 975	
authentication-algorithm statement		
BGP.....	1024	
authentication-key statement		
LDP.....	1026	
usage guidelines.....	630	
RSVP.....	975	
usage guidelines.....	476	
authentication-key-chain statement.....	1027	
auto-bandwidth		
PCE-controlled LSP.....	742	
auto-bandwidth statement.....	838	
usage guidelines.....	258	
auto-policing statement.....	840	
usage guidelines.....	349	
auto-targeted-session statement		
LDP.....	1028	
autobw-state (tracing flag).....	954	
automatic bandwidth allocation.....	257, 838	
bandwidth monitoring.....	263	
LSPs.....	257, 1168	
manually trigger.....	263	
MPLS, statistics.....	264	
threshold.....	260	
automatic mesh, RSVP.....	482	
automatic policers		
LSP bandwidth, changing.....	350	
LSPs.....	350	
overview.....	349	
point-to-multipoint LSPs.....	351	
automatic reoptimization, bypass LSPs.....	507	
B		
backup paths.....	34	
backup-pe-group statement.....	841	
bandwidth		
automatic allocation, LSPs.....	257	
LSP paths.....	256	
RSVP reservations.....	1011	
bandwidth load balancing		
show route protocol rsvp detail command		
.....	809	
verifying	808	
bandwidth model.....	316	
bandwidth oversubscription		
overview.....	319	
bandwidth statement		
fast reroute.....	842	
usage guidelines.....	226	
link protection.....	976	
usage guidelines.....	504	
LSPs		
usage guidelines.....	326	
MPLS		
usage guidelines.....	274	
multiclass LSPs.....	842	
usage guidelines.....	329	
RSVP.....	976	
signaled LSPs.....	842	
usage guidelines.....	256	
static LSPs.....	843	
usage guidelines (ingress router).....	271	
bandwidth update threshold.....	476	
bandwidth, allocating for LSPs.....	1168	
bandwidth-model statement.....	844	
usage guidelines.....	316	
bandwidth-percent statement.....	845	
usage guidelines.....	327, 330	
BFD		
ECMP paths.....	546	
fast reroute.....	355	
LDP LSPs.....	543, 546	
local protection, rapid convergence.....	357	
revert timer.....	215, 933	
RSVP LSPs.....	354, 356	
bfd-liveness-detection statement		
LDP LSPs.....	1029	
usage guidelines.....	543	
RSVP LSPs.....	846	
usage guidelines.....	355	
BGP		
authentication algorithm.....	1024	
destinations.....	40	
link-state distribution.....	8	
BGP link-state distribution.....	10	
BGP-TE NLRI.....	13	
credibility.....	12	
implementation.....	11	
overview.....	10	
supported and unsupported features.....	15	
TLVs.....	13	
traffic engineering database export.....	12	
traffic engineering database import.....	11	
BGP link-state distribution		
configuring.....	364	

- BGP LSPs
 - ping interval.....1163
- BGP-TE NLRI.....13
- BGP-TE TLVs.....13
- binding (tracing flag).....1078
- BMP
 - authentication algorithm.....1024
- braces, in configuration statements.....xl
- brackets
 - angle, in syntax descriptions.....xl
 - square, in configuration statements.....xl
- branch LSPs.....283
- bypass LSPs.....503
 - administrative groups.....504
 - bandwidth.....504
 - bandwidth subscription.....508
 - class-of-service.....505
 - CSPF, disabling.....506
 - explicit paths.....508
 - hop limit.....505
 - maximum number.....506
 - multiple.....496
 - node protection, disabling.....507
 - optimization interval.....507
 - priority and preemption.....509
 - switching away from a network node.....479
 - types.....497
- bypass LSPs, testing.....1266
- bypass statement
 - RSVP.....977
 - static LSP.....978
 - usage guidelines.....503
- C**
- CAC
 - displaying for LSPs.....1188
- call admission control *See* CAC
- Caveats and limitations.....566
- CCC
 - aggregated Ethernet.....650
 - BPDU, nonstandard.....644
 - connections, displaying.....1170, 1356
 - encapsulation
 - Ethernet CCC.....649
 - example configurations.....653, 658
 - graceful restart
 - configuration.....665
 - overview.....664
 - Layer 2 switching cross-connects
 - configuration.....647
 - MPLS tunneling
 - cross-connects.....655, 657, 1100
 - ping CCC LSPs.....360
 - point-to-multipoint LSPs.....665, 1098, 1099
 - traffic policing.....643
- checklists
 - MPLS layered model.....798
- circuit cross-connect *See* CCC
- Cisco HDLC circuits.....648
- Cisco HDLC encapsulation
 - Layer 2 switching cross-connect.....659
- class types
 - bandwidth subscription.....323
- class-of-service statement
 - bypass LSPs.....979
 - usage guidelines.....505
- ingress routers.....847
 - usage guidelines.....271
- signaled LSPs.....847
 - usage guidelines.....245
- static LSPs.....847
- clear ldp neighbor command.....1304
- clear ldp session command.....1305
- clear ldp statistics command.....1306
- clear mpls container-lsp command.....1147
- clear path-computation-client statistics
 - command.....1376
- clear rsdp session command.....1260
- clear rsdp statistics command.....1262
- CLI-controlled LSP.....741
- colored links.....226, 240, 831
- comments, in configuration statements.....xl
- connection (tracing flag).....954
- connection-detail (tracing flag).....954
- connections
 - testing
 - MPLS BGP connections.....1163
 - MPLS LDP connections.....1307
 - MPLS LSP-endpoint connections.....1165
 - MPLS RSVP connections.....1266
- connections statement.....1086
 - complete hierarchy under.....824
- TCC
 - usage guidelines.....663
- Constrained Shortest Path First *See* CSPF
- Constrained Shortest Path First algorithm *See* CSPF algorithm

constrained-path LSPs	
computation	
CSPF algorithm.....	29
disabling.....	239, 903, 951
overview.....	29
example configuration.....	66
overview.....	29
scope.....	29
Constraint-based routing.....	399
Container LSP	
configuration statements.....	403
configuring.....	409, 413
implementation.....	388
implementation overview.....	386
link protection.....	393
LSP merging.....	391
LSP splitting.....	389
network performance impact.....	407
node protection.....	393
normalization.....	394
overview.....	382
protection.....	393
sampling.....	400
supported features.....	408
terminology.....	388
unsupported features.....	408
Container LSPs	
MPLS, displaying.....	1190
container-label-switched-path.....	848
control-channel statement.....	1102
usage guidelines.....	681
conventions	
text and syntax.....	xxxix
corouted-bidirectional statement	
usage guidelines.....	220, 849
corouted-bidirectional-passive statement	
usage guidelines.....	220, 849
CoS.....	316
Differentiated Services.....	324
CoS requests using RSVP.....	459
CoS values.....	244
Credibility.....	12
credibility statement.....	850
cross-connect, circuit See CCC	
CSPF	
statistics, displaying.....	1199
cspf (tracing flag).....	955
CSPF algorithm	
fate-sharing.....	64
offline path computation.....	6, 33
online path computation.....	29
disabling.....	239, 903, 951
overview.....	6
cspf-link (tracing flag).....	955
cspf-node (tracing flag).....	955
curly braces, in configuration statements.....	xi
customer support.....	xli
contacting JTAC.....	xli
D	
damping	
LSP transitions.....	268
database statement.....	851
dead-interval statement.....	1103
usage guidelines.....	686
deaggregate statement.....	1030
usage guidelines.....	541
Delay measurement	
configuring.....	438, 439
definition.....	51
overview.....	50
Delay measurement accuracy.....	52
description statement	
MPLS.....	854
usage guidelines.....	217, 274
static LSPs	
usage guidelines (ingress router).....	271
destination-networks statement	
tunnel.....	980
detail (tracing flag modifier)	
LDP.....	1079
LMP.....	1121
RSVP.....	1013
detours See fast reroute	
devices statement.....	981
Differentiated Services	
bandwidth model.....	316
extended MAM.....	316
interface bandwidth constraints.....	324
LSPs.....	312
MAM.....	316
RDM.....	316
DiffServ	
classes, displaying for MPLS.....	1201
diffserv-te statement.....	856
usage guidelines.....	315

-
- disable (tracing flag modifier).....1079
 - disable option to traceoptions statement
 - LDP.....1078
 - LMP.....1120
 - RSVP.....1012
 - disable statement
 - GMPLS.....1104
 - usage guidelines.....685
 - LDP.....1031
 - usage guidelines.....528
 - link protection.....982
 - usage guidelines.....509
 - MPLS.....857
 - usage guidelines.....239
 - OSPF.....1105
 - usage guidelines.....686
 - RSVP.....982
 - usage guidelines.....471
 - RSVP graceful restart.....982
 - usage guidelines.....514
 - discovery messages, LDP.....521
 - distinct reservations.....466
 - documentation
 - comments on.....xli
 - dod-request-policy statement.....1032
 - downstream on demand, LDP.....592
 - downstream-on-demand statement.....1032
 - usage guidelines.....592
 - DSCP
 - MPLS-tagged packets.....352
 - Dynamic bandwidth management.....409, 413
 - dynamic LSP metric.....230
 - dynamic tunnels.....858
 - destination.....980
 - dynamic-tunnels statement.....858
 - E**
 - ECMP paths
 - BFD.....546
 - ecmp statement.....1033
 - usage guidelines.....546
 - edit protocols mpls command794
 - egress policy, loopback address.....540
 - egress protection
 - Layer 3 VPN
 - with PLR as Protector.....184
 - MPLS, displaying.....1197, 1203
 - service mirroring.....168
 - egress routers
 - example configuration.....276
 - overview.....28
 - signaled LSPs.....212
 - static LSPs.....274, 882
 - egress-policy statement.....1033
 - usage guidelines.....540
 - egress-protection statement
 - MPLS.....859
 - empty paths.....919
 - encapsulation
 - TCC.....659
 - encapsulation statement
 - Layer 2 switching cross-connect
 - usage guidelines.....648
 - logical interfaces.....1087
 - LSP tunnel cross-connect
 - usage guidelines.....656
 - physical interface.....1091
 - TCC
 - usage guidelines.....659
 - encoding-type statement.....860
 - usage guidelines.....688
 - entropy label.....28
 - ingress policy.....219
 - entropy labels, MPLS.....218
 - entropy-label statement.....860
 - usage guidelines.....218
 - error (tracing flag)
 - LDP.....1078
 - MPLS.....955
 - RSVP.....1012
 - ether-pseudowire statement.....868
 - Ethernet extended VLAN TCC, ARP
 - configuration.....662
 - Ethernet TCC
 - ARP configuration.....662
 - ethernet-ccc encapsulation type.....649
 - ethernet-vlan statement.....861
 - event (tracing flag)
 - LDP.....1078
 - RSVP.....1012
 - exclude statement
 - administrative groups
 - usage guidelines.....240
 - fast reroute
 - usage guidelines.....226
 - exclude-srlg
 - usage guidelines.....90, 95, 116

exclude-srlg statement.....	863	fast reroute.....	498
EXP and IP precedence bits.....	353	BFD.....	355
EXP bits.....	26, 244, 246, 345	configuring.....	870
DSCP values.....	352	detours.....	46
rewrite.....	353	for multicast.....	566, 573, 576
EXP rewrite rule.....	246	multiclass LSPs.....	330
expand-loose-hop statement.....	864	overview.....	45, 226
usage guidelines.....	339	path optimization.....	228
experimental bits See EXP bits		path optimization overview.....	50
Explicit Null label.....	25	PFE fast reroute.....	226, 481
Explicit Route object.....	7	soft preemption.....	238
explicit routes.....	6	traffic-engineered LSPs.....	327
explicit senders, RSVP.....	466	fast-reroute statement.....	870
explicit-null statement		RSVP.....	983
LDP.....	1034	usage guidelines.....	228
usage guidelines.....	629	fate-sharing	
MPLS.....	865	CSPF algorithm.....	64
usage guidelines.....	335	example configuration.....	64
RSVP.....	865	overview.....	34
usage guidelines.....	489	signaled LSPs.....	62, 871
usage guidelines.....	336	fate-sharing statement.....	871
explicit-path LSPs		usage guidelines.....	62
computation, disabling.....	239, 903, 951	FEC filters	
configuring.....	278	displaying for LDP.....	1319
example configuration.....	65	fec statement	
overview.....	29	usage guidelines.....	1036
scope.....	29	FECs.....	519, 547
export statement.....	866, 1034	FF (reservation style).....	466
usage guidelines.....	537	filtering received labels.....	535, 1043
extended administrative groups.....	242	fixed-filter reservation style.....	466
extended MAM.....	316, 844	font conventions.....	xxxix
external control.....	741	forwarding See MPLS	
external path computation		forwarding adjacency	
PCE.....	733	configuration.....	725
F		LSP.....	726
facility backup.....	498	OSPF configuration.....	727
show mpls lsp command	761	peer router address.....	726
verifying.....	761, 768	RSVP configuration.....	726
failed LSPs		forwarding equivalence class See FEC	
fast reroute.....	45, 226, 842, 870	forwarding equivalence classes See FECs	
standby secondary paths.....	45	forwarding next hop.....	42
failure-action statement		forwarding table	
LDP LSPs.....	1035	route entries, displaying.....	1360
usage guidelines.....	546	Frame Relay circuits.....	651, 656
RSVP LSPs.....	867	Frame Relay encapsulation	
usage guidelines.....	356	Layer 2 TCC.....	660
family mpls statement.....	868	from statement	
usage guidelines.....	232	MPLS.....	872
		usage guidelines.....	212

FRR	
for multicast.....	566
FRR (fast reroute)	
verifying.....	772
G	
G-Ach.....	52
Generic associated channel.....	52
GMPLS	
Admin Status object.....	690
graceful deletion timeout interval.....	691
graceful LSP teardown.....	690
link-management information, displaying	
all.....	1173
peers.....	1177
routing process.....	1179
statistics.....	1182
traffic-engineered links.....	1184
non-packet LSPs.....	690
permanent LSP deletion.....	691
supported software standards.....	673
temporary LSP deletion.....	690
VLAN LSP.....	698
GMPLS RSVP-TE	
VLAN LSP.....	692
gpip statement.....	873
usage guidelines.....	688
graceful deletion timeout interval.....	691
graceful restart	
LDP.....	1037
point-to-multipoint LSPs.....	305
RSVP.....	985
graceful teardown, GMPLS LSPs.....	690
graceful-deletion-timeout statement.....	984
usage guidelines.....	691
graceful-restart (tracing flag).....	955
graceful-restart statement	
LDP.....	1037
usage guidelines.....	533
RSVP.....	985
usage guidelines.....	514
gre statement.....	874
GRE tunnels.....	70
groups	
administrative.....	226, 240, 831
MPLS, displaying administrative.....	1186
H	
headers, MPLS and IPv4.....	353
hello acknowledgments	
RSVP.....	479
hello interval	
LDP.....	528, 1038
RSVP.....	475, 986
hello messages.....	521
interface.....	475
node ID.....	478
hello packets	
RSVP.....	462
hello-acknowledgements statement.....	986
RSVP	
usage guidelines.....	479
hello-dead-interval statement.....	1106
usage guidelines.....	683
hello-interval statement	
LDP.....	1038
usage guidelines.....	528
LMP.....	1107
usage guidelines.....	683
OSPF.....	1108
usage guidelines.....	686
RSVP.....	986
usage guidelines.....	475
hello-packets (tracing flag)	
LMP.....	1120
helper-disable statement	
LDP.....	1039
usage guidelines.....	533
RSVP	
usage guidelines.....	514
hold priority.....	250
hold time	
LDP.....	529, 1041
signaled LSPs.....	268, 835
hold-time statement	
LDP.....	1041
usage guidelines.....	529
holddown-interval statement.....	1040
usage guidelines.....	547
hop-limit statement.....	875, 987
usage guidelines.....	256, 505
host routes.....	40, 228
hosts, reachability	
MPLS BGP LSPs.....	1163
MPLS LDP LSPs.....	1307
MPLS LSP endpoints.....	1165
MPLS RSVP LSPs.....	1266
hot-standby state.....	267

I	
IEEE 802.p rewrite rule.....	246
ignore-lsp-metrics statement.....	1042
usage guidelines.....	630
IGP	
advertising LSPs.....	38
destinations.....	41
limitations.....	9
role.....	8
shortcuts	
enabling.....	36
LSP metrics.....	231
overview.....	34
qualified LSPs.....	36
routing tables.....	37
uses of.....	36
IGP synchronization, LDP.....	633
igp-synchronization statement.....	1042
usage guidelines.....	633
Implicit Null label.....	25
import statement.....	876
LDP.....	1043
usage guidelines.....	535
in-band signaling	
for multipoint LDP.....	607
inband signaling	
for multipoint LDP.....	598
include statement	
fast reroute	
usage guidelines.....	226
include-all statement	
administrative groups	
usage guidelines.....	240
include-any statement	
administrative groups	
usage guidelines.....	240
inet.0 routing table	
IGP shortcuts.....	37
MPLS.....	43
inet.3 or inet6.3 routing table	
routes, installing.....	228
inet.3 routing table	
IGP shortcuts.....	37
MPLS.....	43
information distribution, traffic engineering.....	5
ingress routers	
address configuration.....	212, 872
configuring for static LSPs.....	271
example configurations.....	273
overview.....	28
path connection retry information.....	229, 932
ingress statement	
static LSP.....	879
ingress static LSPs.....	271
ingress-policy statement.....	881, 1044
usage guidelines.....	219, 547
init (tracing flag)	
LMP.....	1120
initialization (tracing flag).....	1079
install statement	
MPLS.....	880
usage guidelines.....	228
static LSPs	
usage guidelines (ingress router).....	271
Integrity object.....	461
inter-area traffic engineering.....	339
Inter-domain point-to-multipoint LSPs.....	303
inter-domain statement.....	225, 883
interface (from operator, LDP).....	535
interface statement	
LDP.....	1045
usage guidelines.....	528
LMP.....	1109
usage guidelines.....	677
multicast.....	1076
RSVP.....	988
usage guidelines.....	471
static LSPs.....	882
interface-switch statement.....	1096
Layer 2 switching cross-connects	
usage guidelines.....	652
usage guidelines.....	652
interfaces	
aggregated.....	39
configuration, incorrect for MPLS.....	760
MPLS, verifying	759
interior gateway protocol See IGP	
intermediate routers	
configuring for static LSPs.....	274, 882
example configurations.....	275
Internet draft	
draft-ietf-mpls-rsvp-te-no-php-oob-mapping-01.txt, Non PHP behavior and Out-of-Band Mapping for RSVP-TE LSPs.....	22
Internet draft	
draft-napierala-mpls-targeted-mldp-01.txt, Using LDP Multipoint Extensions on Targeted LDP Sessions	520

- intraregion LSPs.....36
- IP packets over aggregated interfaces.....39
- IPv4 Explicit Null label.....25
- IPv6
 - Implicit Null label.....25
 - tunneling over MPLS.....71
- ipv6-tunneling statement.....883
- IS-IS
 - authentication, displaying.....1349
- K**
- keep multiplier, RSVP.....483, 989
- keep-multiplier statement.....989
 - usage guidelines.....483
- keepalive-interval statement.....1046
 - usage guidelines.....531
- keepalive-timeout statement.....1047
 - usage guidelines.....531
- keepalives
 - interval.....531, 1046
 - timeout.....531, 1047
- keyboard sequences
 - used with monitor mpls command.....1263
- L**
- l2-smart-policy statement.....1047
 - usage guidelines.....542
- label (tracing flag).....1079
- Label Distribution Protocol See LDP
- label filtering.....535, 1043
- Label object.....7
- Label operation.....557
- Label Request object.....7
- label-switched paths See LSPs
- label-switched-path statement
 - GMPLS.....1109
 - MPLS.....884
 - MPLS with RSVP.....884
 - usage guidelines.....472
 - usage guidelines.....726
- label-switched-path-template.....887
- label-switched-path-template statement.....990
- label-withdrawal-delay statement.....1048
 - usage guidelines.....633
- labeled-unicast statement
 - usage guidelines.....336
- labels
 - allocation.....26
 - numerical ranges.....24
 - operations.....27, 523
 - overview.....20
 - properties.....271
 - reserved labels.....25
 - stacks.....26
 - swapping.....4
 - values.....24
- Layer 2 switching
 - MPLS.....663
 - TCC.....663
- Layer 2 switching cross-connect
 - CCC connections.....652
 - CCC encapsulation.....648
 - configuration.....647
 - configuring MPLS.....652
 - example configuration.....653
 - TCC encapsulation.....659
- Layer 2 VPNs
 - aggregated Ethernet.....650
- Layer 3 VPN
 - egress protection
 - with PLR as Protector.....184
- layered model
 - checklist798
 - figure799
 - MPLS layer, figure784
- LDP
 - authentication algorithm.....1024
 - authentication keychain.....1027
 - BFD.....543, 546
 - carrier-of-carriers VPNs.....629
 - configuring.....1045, 1049
 - database entries, displaying.....1310
 - disabling.....528, 1031
 - downstream on demand.....592
 - ECMP-aware BFD.....546
 - egress policy.....540
 - enabling.....528
 - example configuration
 - received label filtering.....536
 - tracing.....639
 - Explicit Null label.....489, 629
 - FEC filters, displaying.....1319
 - FEC policers.....541
 - graceful restart.....532, 533, 1037
 - hello interval.....528, 1038
 - hello messages.....521
 - hold time.....529, 1041
 - IGP synchronization.....633

Implicit Null label.....	489, 629
interfaces, displaying.....	1320
introduction.....	550
Junos implementation.....	520
Junos OS implementation.....	550
keepalive	
interval.....	531, 1046
timeout.....	531, 1047
label operations.....	523
link protection.....	549
caveats and limitations.....	566
configuring.....	548
label operation.....	557
limitations.....	552
modes of operation.....	555
LSPs, displaying.....	1329
message types.....	521
metrics.....	628
minimum configuration.....	528
multiple instances.....	629
multipoint.....	598, 607
multipoint extensions.....	550
neighbors	
clearing connections.....	1304
displaying.....	1322
OAM ingress policy.....	547
OAM periodic traceroute.....	634
operations.....	521
overview.....	519
policy filters.....	1043
received label filtering.....	535, 1043
route preferences.....	532, 1071
routes, displaying.....	1331
session protection	
configuration.....	631
overview.....	525
sessions	
clearing.....	1305
displaying.....	1335
statistics	
clearing.....	1306
displaying.....	1341
supported software standards.....	520
synchronization with the IGP.....	632
targeted hello messages.....	521
teardown-delay	
maximum-sessions.....	1028
timer.....	528, 1038
tracing LSPs.....	1352
tracing operation of.....	637, 1078
traffic statistics, displaying.....	1345
tunneling through RSVP LSPs.....	523, 629, 887
ultimate-hop popping.....	489, 629, 1034
LDP LSPs	
ping interval.....	1307
ldp statement.....	1049
complete hierarchy under.....	1021
usage guidelines.....	528
ldp-synchronization statement.....	1052
usage guidelines.....	632
ldp-tunneling statement.....	887
usage guidelines.....	629
least-fill statement.....	930
usage guidelines.....	232
least-fill tie-breaking rule.....	32, 232, 930
link attributes considered by CSPF algorithm.....	29
link coloring.....	226, 240, 831
link hello messages, LDP.....	1038
link protection.....	498, 509
bypass LSPs	
administrative groups.....	504
many-to-one backup	
verifying	768
multiple bypass LSPs.....	506
RSVP.....	495
show mpls lsp extensive command.....	768
show rsvp interface command.....	772
show rsvp session detail command.....	770
soft preemption.....	238
static LSPs.....	509
switching away from a network node.....	479
Link protection.....	393
Link state distribution	
BGP.....	8
link-layer protocols.....	23
link-management statement.....	1110
complete hierarchy under.....	824
usage guidelines.....	677
link-node protection	
between autonomous systems.....	150
link-protection statement	
MPLS	
usage guidelines.....	274, 304, 509
RSVP.....	992
usage guidelines.....	502
signaled LSPs.....	888
static LSPs.....	889
usage guidelines (ingress router).....	271

-
- Link-state distribution
 - need.....9
 - using BGP.....362, 364
 - link-state paths *See* LSPs
 - LMP
 - peer network device configuration.....680
 - tracing protocol operations.....1120
 - tracing protocol traffic.....687
 - traffic engineering links.....677
 - lmp (tracing flag).....1012
 - lmp-control-channel statement.....1111
 - usage guidelines.....681
 - lmp-protocol statement.....1111
 - load balancing
 - MPLS.....218
 - MPLS LSPs.....232
 - verifying.....805
 - load-balance statement.....993
 - usage guidelines.....481
 - load-balance-label-capability statement.....889
 - local control.....741
 - local protection, BFD.....357
 - local-address statement
 - link management.....1112
 - usage guidelines.....679
 - usage guidelines.....725
 - log-updown statement
 - LDP.....1053
 - usage guidelines.....632
 - MPLS.....890
 - usage guidelines.....344
 - logical-router *See* logical-system
 - logical-routers *See* logical-systems
 - loopback address, egress policy.....540
 - loose explicit routes.....6, 278
 - Loss and delay measurement
 - concepts.....52
 - definition.....51
 - functionality.....55
 - importance.....50
 - mechanisms.....51
 - metrics.....52
 - supported and unsupported features.....56
 - Loss measurement
 - configuring.....438, 439
 - definition.....51
 - overview.....50
 - Loss measurement mode.....52
 - Loss measurement synchronization.....52
 - LSP
 - behavior.....740
 - control mode.....741
 - route, checking.....788
 - types.....741
 - LSP configuration
 - support for PCE-controlled LSP.....742
 - LSP graceful teardown.....690
 - LSP merging
 - constraints.....392
 - overview.....392
 - triggers.....392
 - LSP metric.....339
 - LSP protection
 - PCE-controlled LSP.....742
 - LSP splitting
 - constraints.....390
 - overview.....389
 - supported criteria.....390
 - triggers.....391
 - lsp-attributes statement.....894
 - usage guidelines.....688
 - lsp-history (tracing flag).....955
 - lsp-next-hop, static routes.....918
 - lsp-ping-interval statement
 - LDP LSPs.....1062
 - RSVP LSPs.....911
 - lsp-switch statement.....1097
 - lsping (tracing flag).....955
 - LSPs
 - adaptive rerouting.....249, 825
 - administrative groups
 - admin-groups statement.....831
 - configuring.....240
 - fast reroute.....226
 - advertising in IGPs.....38
 - associating addresses.....228, 880
 - attributes considered by CSPF algorithm.....29
 - automatic bandwidth.....1227
 - automatic bandwidth allocation.....838
 - automatic policers.....350
 - bandwidth
 - maximum bounds.....259
 - minimum bounds.....259
 - bandwidth allocation, adjusting.....1168
 - BFD configuration.....354
 - BGP, ping interval.....1163
 - bypass.....509
 - CAC information, displaying.....1188

clearing.....	1145	priorities.....	250, 927
configuration statements.....	884	recording routes.....	244
configuring.....	326	reoptimization.....	251, 913, 916
constrained-path See constrained-path LSPs		router functions.....	28
CoS values.....	244	routing options.....	7
damping LSP transitions.....	268	RSVP, ping interval.....	1266
description, textual.....	217	RSVP, real-time status.....	1263
differentiated service aware.....	312	RSVP-signaled.....	29
egress routers.....	212, 274, 276, 882	scope of.....	29
entropy labels.....	218	secondary.....	214, 937
explicit-path See explicit-path LSPs		signaled See signaled LSPs	
failure of.....	45	soft preemption.....	238
fast reroute.....	45, 226, 842, 870	standby secondary paths.....	45
fate-sharing.....	34, 62, 871	standby state.....	267, 944
forwarding next hops		static See static LSPs	
selecting.....	42	switching away from a network node.....	479
hold time.....	268, 835	text description.....	854
hop limit.....	256	tie-breaking rules.....	32, 232, 930
host routes.....	40	traffic engineering, configuring.....	336
IGP shortcuts.....	34	TTL decrementing, disabling.....	236, 904, 907
ingress routers.....	212, 872	tunnel cross-connects, MTU.....	655
inter-domain.....	225	tunneling through RSVP LSPs.....	523, 629, 887
intermediate routers.....	274, 882	ultimate-hop popping.....	222, 486
intraregion LSPs.....	36		
LDP, displaying.....	1329	M	
LDP, ping interval.....	1307	make-before-break.....	246
load balancing.....	232	Make-before-break.....	564
MBB switchover.....	246	MAM.....	316, 844
metrics.....	230, 231, 898	manuals	
MPLS routers, configuring.....	68	comments on.....	xli
MPLS, displaying.....	1209	many-to-one backup	
multiple bypass.....	496	show mpls lsp command	761
named paths.....	60, 919	verifying	761, 768
OAM configuration.....	354	max-bypasses statement.....	994
overview.....	4, 24	usage guidelines.....	506
packet traversal.....	5, 28	maximum-bandwidth statement.....	894
path		usage guidelines.....	259
bandwidth.....	256	maximum-labels statement.....	895
calculation.....	3	maximum-neighbor-recovery-time	
connection retry information.....	229, 932	statement.....	1055
length.....	256, 875, 987	usage guidelines.....	535
smart optimize timer.....	255	maximum-recovery-time statement.....	1055
pings.....	359	MD5 authentication.....	476
ping interval, LDP.....	545	Measurement point.....	52
ping interval, RSVP.....	356	messages	
policing.....	347	LDP message types.....	521
preemption.....	250, 927	MPLS system log.....	344, 890
preference levels.....	243, 925	Resv, RSVP.....	465
primary.....	214, 926	ResvConfirm, RSVP.....	466

ResvErr, RSVP.....	466	BGP destinations.....	40
ResvTear, RSVP.....	465	BGP-signaled LSP connections	
RSVP message types.....	463	operability, checking.....	1163
RSVP PathErr.....	68	CCC connections, displaying.....	1170, 1356
RSVP refresh.....	483	configuring.....	59
metric statement		CoS values.....	244
MPLS.....	898	CSPF statistics, displaying.....	1199
usage guidelines.....	231	DiffServ classes, displaying.....	1201
static LSPs		DSCP and EXP values.....	352
usage guidelines (ingress router).....	271	entropy label, configuration.....	218
metrics		EXP bits.....	26, 244, 246, 345
dynamic LSP metric.....	230	Explicit Null label.....	335
LDP tracking IGP.....	628	extended administrative groups.....	242
static LSP metric.....	231, 898	fast reroute.....	45, 226, 842, 870
minimum-bandwidth statement.....	898	firewall filter.....	345
usage guidelines.....	259	GRE tunnels.....	70
minimum-bandwidth-adjust-interval		IGP and BGP destinations.....	41
statement.....	896	Implicit Null label.....	335
minimum-bandwidth-adjust-threshold-change		inter-AS link-node protection.....	150
statement.....	896	interfaces, displaying.....	1205
minimum-bandwidth-adjust-threshold-value		IPv4 packet headers.....	353
statement.....	897	IPv6.....	71
mldp-inband-signalling statement.....	1056	label range.....	24
model, checklist for.....	798	Layer 2 switching TCC.....	663
MoFRR		LDP-signaled LSP connections	
multipoint LDP.....	573	operability, checking.....	1307
PIM.....	573, 576	link-layer protocols supported.....	23
mofrr-asm-starg statement.....	1057	link-management information, displaying	
mofrr-disjoint-upstream-only statement.....	1058	all.....	1173
mofrr-no-backup-join statement.....	1059	peers.....	1177
mofrr-primary-selection-by-routing		routing process.....	1179
statement.....	1060	statistics.....	1182
monitor label-switched-path command.....	1263	traffic-engineered links.....	1184
monitor mpls delay rsvp command.....	1150	load balancing.....	218, 232
monitor mpls loss rsvp command.....	1154	LSP endpoint connections	
monitor mpls loss-delay rsvp command.....	1159	operability, checking.....	1165
monitor-bandwidth statement.....	899	LSP tunnel cross-connects	
usage guidelines.....	263	MTU.....	655
most-fill statement.....	930	LSPs See LSPs	
usage guidelines.....	232	LSPs, displaying.....	1230
most-fill tie-breaking rule.....	32, 232, 930	OAM.....	354
MPLS.....	3	overview.....	20
administrative groups, displaying.....	1186	packets over aggregated interfaces.....	39
aggregated interfaces.....	39	ping	
automatic bandwidth allocation.....	838	Layer 3 VPNs.....	360
automatic bandwidth allocation,		LSP end points.....	360
statistics.....	264	LSPs.....	359
backbones, packet traversal.....	5, 28	routing tables.....	43
BFD.....	354, 356	RSVP See RSVP	

signaled LSPs See signaled LSPs	
smart optimize timer.....	255
SNMP traps.....	344, 890
soft preemption.....	238
standby secondary paths.....	45
static.....	271, 882
static LSPs See static LSPs	
static LSPs, displaying.....	1233
statistics output.....	342
supported software standards.....	20
system log messages.....	344, 890
tracing LSPs.....	1298, 1352
tracing protocol operations.....	361, 954
traffic engineering.....	339
overview.....	24
traffic protection.....	45
traffic statistics.....	342, 947
traffic-engineering	
database.....	851, 866, 876 See export See import
tunneling	
CCC connection.....	657, 1100
CCC encapsulation.....	656
example configurations.....	658
IPv6.....	71
overview.....	655
ultimate-hop popping.....	222, 335, 486, 865
See also LDP, LSPs, RSVP, traffic engineering	
database	
MPLS configuration	
support for PCE-controlled LSP.....	742
MPLS layer	
broken network topology, figure	784
checking	783
MPLS protocol	
activate interface command.....	794
administrative groups.....	760
configuration, incorrect.....	761
edit protocols mpls command.....	794
interfaces, verifying.....	759
labels, verifying.....	791
ping command.....	791
ping mpls rsvp lsp-name detail command.....	791
routing table	788
show configuration interfaces command.....	793
show configuration protocols mpls	
command.....	792
show mpls interface command.....	759
show mpls lsp extensive command.....	785, 795
show route command.....	789
show route table mpls.0 command.....	788
MPLS RSVP-TE	
understanding.....	733
mpls statement	
Layer 2 switching cross-connect.....	899
usage guidelines.....	652
Layer 2 switching TCC.....	663
MPLS.....	899
complete hierarchy under.....	817
usage guidelines.....	59
mpls transport profile oam	
overview.....	136
mpls-tp-mode	
usage guidelines.....	136
mpls-tp-mode statement	
MPLS-TP.....	900
mpls.0 routing table.....	43
MTU signaling, in RSVP.....	468
mtu-signaling statement.....	900
usage guidelines.....	484
multicast	
fast reroute.....	566
RPF check policy.....	306
Multicast LDP	
extension on targeted LDP session.....	551
link protection.....	555
make-before-break.....	564
sample link protection configuration.....	563
Multicast-only fast reroute.....	573, 576
multiclass LSPs	
bandwidth subscription.....	323
configuring.....	329
fast reroute.....	330
multiple bypass LSPs.....	496, 503, 504, 506
multiple push (label operation).....	27
multipoint	
LDP.....	598, 607
multipoint LDP	
MoFRR.....	573
N	
named paths	
empty paths.....	919
example configuration.....	62
overview.....	60
neighbor (from operator, LDP).....	535
network	
problems.....	805

next hop (from operator, LDP).....	535
next hops	
selecting.....	42
next-hop bypass LSP.....	497
next-hop statement.....	901
MPLS	
usage guidelines.....	274
static LSPs	
usage guidelines (ingress router).....	271
next-next-hop bypass LSP.....	497
no-adjacency-down-notification statement	
IS-IS.....	997
no-aggregate statement.....	974
usage guidelines.....	473
no-bfd-triggered-local-repair statement.....	902
no-cspf statement.....	903, 951, 998
usage guidelines.....	239, 506
no-decrement-ttl statement.....	904
no-forwarding statement.....	1061
usage guidelines.....	628
no-install-to-address statement.....	905
static LSPs	
usage guidelines (ingress router).....	271
usage guidelines.....	213
no-interface-hello statement.....	998
RSVP	
usage guidelines.....	478
no-load-balance-label-capability statement.....	905
no-local-reversion statement.....	995
no-mcast-replication statement.....	906
no-neighbor-down-notification statement.....	999
no-node-id-subobject statement.....	999
usage guidelines.....	510
no-node-protection statement	
usage guidelines.....	507
no-p2mp-sublsp statement.....	1000
usage guidelines.....	308
no-propagate-ttl statement.....	907
no-record statement.....	931
usage guidelines.....	244
no-reliable statement.....	1006
usage guidelines.....	473
no-trap statement.....	908
usage guidelines.....	344
no-world-readable option to traceoptions	
statement	
LDP.....	1079
LMP.....	1121
MPLS.....	955
RSVP.....	1013
node ID hellos, RSVP.....	478
node protection.....	497, 498
soft preemption.....	238
static LSPs.....	509
switching away from a network node.....	479
Node protection.....	393
node-hello statement.....	996
RSVP	
usage guidelines.....	478
node-link protection	
show mpls lsp command	761
show mpls lsp extensive.....	763
show rsvp interface command.....	764
show rsvp interface extensive command.....	765
show rsvp session detail command.....	766
verifying	761
node-link-protection statement.....	1000
usage guidelines.....	509
node-protection statement	
MPLS	
usage guidelines.....	274
static LSPs.....	909
usage guidelines (ingress router).....	271
Normalization	
auto-bandwidth inter-operation.....	395
computed range.....	397
constraints.....	394
feasible range.....	397
overview.....	394
normalization.....	910
notification (tracing flag).....	1079
notification messages	
LDP.....	522
nsr-synchronization (tracing flag).....	955
nsr-synchronization-detail (tracing flag).....	955
O	
OAM.....	52
ingress policy for LDP LSPs.....	547
OAM periodic traceroute, LDP.....	634
oam statement	
LDP LSPs.....	1062
usage guidelines.....	543
RSVP LSPs.....	911
usage guidelines.....	355
OAM, MPLS.....	354
offline path calculation.....	6, 33

On-demand loss and delay measurement		packet headers, MPLS and IPv4.....	353
overview.....	50	packet traversal on LSPs.....	5, 28
On-Demand Loss and Delay		packets (tracing flag)	
Measurement.....	438, 439	LDP.....	1079
one-to-one backup.....	498	LMP.....	1120
verifying	772	RSVP.....	1012
operations on labels.....	27	parentheses, in syntax descriptions.....	xl
optimize-adaptive-teardown statement.....	912	parser (tracing flag)	
optimize-aggressive statement.....	913	LMP.....	1120
usage guidelines.....	251	passive statement.....	1112
optimize-hold-dead-delay statement		usage guidelines.....	684
usage guidelines.....	246, 251	path	
optimize-switchover-delay statement		bandwidth, LSP.....	256
usage guidelines.....	246, 251	calculation	
optimize-timer statement		constrained-path	
bypass LSPs.....	1001	computation.....	239, 903, 951
usage guidelines.....	507	CSPF algorithm.....	6, 29
MPLS.....	916	offline path computation.....	6, 33
usage guidelines.....	251	routing options.....	7
optimizing LSPs.....	246, 251, 913, 916	tie-breaking rules.....	32, 232, 930
OSPF		connection retry information.....	229, 932
hello interval.....	1108	length, LSP.....	256, 875, 987
inter-area traffic engineering.....	339	selection component, traffic engineering.....	6
link-state		path (tracing flag)	
advertisements.....	1117	LDP.....	1079
LSP metric advertisement.....	339	RSVP.....	1012
peer interfaces.....	1114	path computation	
router dead interval.....	1103	PCE.....	731
transmission delay.....	1122	<i>See also</i> PCC	
outgoing MTU value in RSVP		Path Computation Client.....	738
determining.....	469	Path Computation Element.....	737
output control keys		Path Computation Element Protocol.....	738
for monitor mpls command.....	1263	path messages, RSVP.....	465
P		path optimization	
P2MP LSPs, testing.....	1266	fast reroute.....	228
P2MP LSPs, tracing.....	1298	path selection.....	216
p2mp statement.....	917, 1063	path statement	
usage guidelines.....	283	MPLS.....	919
p2mp-lsp-next-hop statement.....	918	RSVP.....	1002
RSVP		usage guidelines.....	508
usage guidelines.....	285	path-mtu statement.....	484, 920
usage guidelines.....	277	PathErr messages.....	68, 466
p2mp-receive-switch statement.....	1098	pathtear (tracing flag).....	1012
usage guidelines.....	667	PathTear messages, RSVP.....	465
p2mp-transmit-switch statement.....	1099		
usage guidelines.....	666		
packet forwarding component			
traffic engineering.....	4		

PCC.....	738	traffic-class (loss-delay).....	960
active pce.....	1378	traffic-class(loss).....	958
active-PCE		periodic (tracing flag).....	1079
request.....	1377	periodic-traceroute statement.....	1065
statistics.....	1382	usage guidelines.....	546, 634
clear.....	1376	permanent GMPLS LSP deletion.....	691
PCE.....	737	PFE fast reroute.....	226, 481
external path computing.....	743	PIM	
impact on network performance.....	743	mldp-inband-signalling statement.....	1056
path computation.....	733	MoFRR.....	573, 576
PCE and PCC interaction.....	739	policy statement.....	1070
PCE-controlled LSP.....	741	ping	
auto-bandwidth.....	742	Layer 3 VPNs.....	360
LSP protection.....	742	LSP end point.....	360
PCEP.....	731, 733, 738	LSPs.....	359
for MPLS RSVP-TE.....	733	point-to-multipoint LSP.....	360
functions.....	738	ping command	791
pcep		ping mpls bgp command.....	1163
configuring.....	743	ping mpls ldp command.....	1307
for MPLS RSVP-TE.....	743	ping mpls lsp-end-point command.....	1165
PCEP session.....	739	ping mpls rsvp command.....	1266
pcep statement		ping mpls rsvp lsp-name detail command	791
PCEP		PLP bit.....	245
complete hierarchy under.....	1125	point-to-multipoint LSPs	
peer network device configuration.....	680	automatic policers.....	351
peer statement		branch LSPs.....	283
LMP.....	1113	dynamic.....	284
usage guidelines.....	680	static.....	284
peer-interface statement.....	1114	CCC.....	665
OSPF		configuration.....	283
usage guidelines.....	685	graceful restart.....	282, 305
RSVP.....	1003	GRES.....	282
usage guidelines.....	685	inter-domain.....	303
usage guidelines.....	726, 727	link protection.....	304
Per-LSP autobandwidth adjustments.....	395	overview.....	281
per-prefix-label statement.....	921	RPF check policy.....	306
performance monitoring		static routes.....	277
configuring.....	448, 450	ultimate-hop popping.....	490
performance-monitoring		with RSVP signaling.....	285
delay (querier).....	852	policers	
delay (responder).....	853	LDP FECs.....	541
loss (querier).....	891	policing.....	347, 349, 643
loss (responder).....	892	policing filter statement	
loss-delay (querier).....	893	usage guidelines.....	347
LSPs, clearing.....	1149	policing statement.....	923, 1067
MPLS.....	922	static LSPs	
querier.....	1136	usage guidelines (ingress router).....	271
traffic-class (delay).....	956	usage guidelines.....	349, 541
		policy filters, LDP.....	1043

policy statement	
for multipoint LDP.....	1070
multicast-only fast reroute.....	1068
pop (label operation).....	27
pop statement	
MPLS.....	924
usage guidelines.....	274
PPP circuits	
Layer 2 switching cross-connects.....	648
preemption	
LSPs.....	238
RSVP sessions.....	484
signaled LSPs.....	250, 927
preemption statement.....	1004
usage guidelines.....	484
preference levels	
LDP routes.....	532, 1071
signaled LSPs.....	243, 925
static LSPs.....	271
preference statement	
LDP.....	1071
usage guidelines.....	532
signaled LSPs.....	925
usage guidelines.....	243
static LSPs.....	925
usage guidelines (ingress router).....	271
primary LSPs.....	214, 926
primary path	
show mpls lsp extensive ingress command	
.....	780
show rsvp interface command.....	780
verifying.....	779
primary paths	
revert timer.....	215
revert timer, BFD.....	933
selection.....	216
primary statement	
MPLS.....	926
usage guidelines.....	214
priorities	
signaled LSPs.....	250, 927
priority statement	
MPLS.....	927
usage guidelines.....	250
RSVP.....	1005
usage guidelines.....	509
Proactive Loss and Delay Measurement.....	448, 450
process (tracing flag).....	1120
protection statement.....	150
protection-revert-time statement.....	928
push (label operation).....	27
push statement	
MPLS.....	929
static LSPs	
usage guidelines (ingress router).....	271
Q	
Querier.....	52
Query rate.....	52
R	
R1 router	
show mpls lsp extensive command	
link protection	768
node-link protection.....	763
show mpls lsp extensive ingress command	
primary path.....	780
show rsvp interface command	
link protection	772
node-link protection	764
primary path	780
show rsvp interface extensive command	
node-link protection	765
show rsvp session detail command	
link protection	770
node-link protection	766
random statement.....	930
usage guidelines.....	232
random tie-breaking rule.....	32, 232, 930
RDM.....	316, 844
real-time monitoring	
RSVP LSPs.....	1263
receive (tracing flag modifier)	
LDP.....	1079
LMP.....	1121
RSVP.....	1013
received label filtering.....	1043
reconnect-time statement.....	1072
usage guidelines.....	534
Record Route object.....	244
record statement.....	931
usage guidelines.....	244
recording routes.....	244
recovery-time statement.....	1073
usage guidelines.....	535
refresh messages, RSVP.....	483
refresh reduction, RSVP.....	467
refresh time, RSVP.....	483

refresh-time statement.....	1006	revert-timer statement.....	933
usage guidelines.....	483	usage guidelines.....	215
regular expressions		rewrite rules.....	246
LSPs, clearing.....	1145	IEEE 802.p and MPLS CoS.....	246
reliable statement.....	1006	MPLS and VPNs.....	353
usage guidelines.....	473	RFC 5317, Joint Working Team (JWT) Report on MPLS Architectural Considerations for a Transport Profile	21
remote-address statement		RFC 5586, MPLS Generic Associated Channel.....	21
control channel management		RFC 5654, Requirements of an MPLS Transport Profile.....	21
usage guidelines.....	682	RFC 5712, MPLS Traffic Engineering Soft Preemption.....	21
LMP control channel.....	1115	RFC 5718, An In-Band Data Communication Network For the MPLS Transport Profile.....	21
LMP traffic engineering.....	1115	RFC 5860, Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks.....	21
usage guidelines.....	679, 726	RFC 5921, A Framework for MPLS in Transport Networks.....	21
remote-id statement		RFC 5950, Network Management Framework for MPLS-based Transport Networks.....	21
link management.....	1116	RFC 5951, Network Management Requirements for MPLS-based Transport Networks.....	21
usage guidelines.....	680	RFC 5960, MPLS Transport Profile Data Plane Architecture.....	21
remote-interface-switch statement.....	1100	RFC 6215, MPLS Transport Profile User-to-Network and Network-to-Network Interfaces.....	22
usage guidelines.....	657	RFC 6291, Guidelines for the Use of the "OAM" Acronym in the IETF.....	22
reoptimizing LSPs.....	251, 913, 916	RFC 6370, MPLS Transport Profile (MPLS-TP) Identifiers	22
request mpls container-lsp command.....	1167	RFC 6371, Operations, Administration, and Maintenance Framework for MPLS-Based Transport Networks.....	22
request mpls lsp adjust-autobandwidth		RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths	22
command.....	1168	RFC 6424, Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels.....	22
request path-computation-client active-pce		RFC 6425, Detecting Data-Plane Failures in Point-to-Multipoint MPLS - Extensions to LSP Ping	22
command.....	1377	RFC 6426, MPLS On-demand Connectivity Verification and Route Tracing.....	22
requests, CoS.....	459		
rerouting LSPs			
adaptive rerouting.....	249, 825		
fast reroute.....	45, 226, 842, 870		
reservation styles.....	466		
reserved labels.....	25		
reserving network resources See RSVP			
resource classes.....	226, 240		
Resource Reservation Protocol See RSVP			
Responder.....	52		
resv (tracing flag).....	1013		
Resv messages, RSVP.....	465		
ResvConfirm messages, RSVP.....	466		
ResvErr messages, RSVP.....	466		
resvtear (tracing flag).....	1013		
ResvTear messages, RSVP.....	465		
retransmission-interval statement.....	1116		
usage guidelines.....	684		
retransmit-interval statement.....	1117		
usage guidelines.....	686		
retry information.....	229, 932		
retry-limit statement.....	932, 1118		
usage guidelines.....	229, 684		
retry-timer statement.....	932		
usage guidelines.....	229		

RFC 6428, Proactive Connectivity Verification, Continuity Check and Remote Defect Indication for MPLS Transport Profile	22	enabling.....	471
route (tracing flag).....	1013	example configurations.....	472, 492
LDP.....	1079	Explicit Null label.....	489
route preferences		for point-to-multipoint LSPs.....	285
LDP.....	532, 1071	graceful restart.....	511, 985
signaled LSPs.....	243, 925	hello acknowledgments.....	479
route-socket (tracing flag)		hello interval.....	475, 986
LMP.....	1120	hello packets.....	462
Router Alert label.....	25	IGP hello packets.....	462
routers		Implicit Null label.....	489
egress See egress routers		interfaces, displaying.....	1271
ingress See ingress routers		Junos implementation.....	461
label operations.....	27	keep multiplier.....	989
LSP functions.....	28	link protection.....	509
transit.....	28	load balancing.....	481
routes		LSP connections	
recording.....	244	operability, checking.....	1266
route preferences.....	243, 532, 925, 1071	LSPs, real-time status.....	1263
routes, displaying		message types.....	463
CCC.....	1359	MPLS, configuring with RSVP.....	472
in the forwarding table.....	1360	MTU signaling in.....	468
in the LDP internal topology table.....	1331	multipath extensions.....	382
routing (tracing flag).....	1120	multipath implementation.....	383
routing options, traffic engineering.....	7	neighbors, displaying.....	1276
routing table, MPLS.....	788	node ID hellos.....	478
routing tables		overview.....	459
host routes, installing.....	228, 880	PathErr messages.....	68
IGP shortcuts.....	37	preemption.....	484
inet.0.....	37, 43	reservation styles.....	466
inet.3.....	37, 43	sessions.....	472
inet.3 or inet6.3.....	228	clearing.....	1260
MPLS.....	43	displaying.....	1281
mpls.0.....	43	setup protection.....	480
rpf-check-policy statement.....	934	signaled LSPs.....	29
configuration guidelines.....	306	signaling extensions.....	7
RRO node ID sub-object, disabling.....	510	statistics	
RSVP.....	3	clearing.....	1262
aggregation.....	974	displaying.....	1291
authentication.....	476, 975	supported software standards.....	460
automatic mesh, configuration.....	482	timers.....	483, 1006
bandwidth		timers, hello packets.....	462
reserving.....	1011	tracing LSPs.....	1298
update threshold.....	476	tracing protocol traffic.....	491, 1012
BFD.....	354, 356	tunneling LDP LSPs through RSVP	
configuration, minimum.....	471	LSPs.....	523, 629, 887
Differentiated Services.....	324	ultimate-hop popping.....	489, 490
disabling.....	471	unnumbered interfaces.....	477
		version, displaying.....	1295
		See also LDP	

RSVP LSP hierarchy	
configuration.....	724
overview.....	723
RSVP LSPs	
ping interval.....	1266
RSVP refresh reduction	
configuration.....	473
overview.....	467
rsvp statement.....	1007
complete hierarchy under.....	971
mpls.....	1007
usage guidelines.....	471
rsvp-error-hold-time statement.....	935
usage guidelines.....	68
RSVP-TE	
addressing current limitations.....	736
current limitations.....	735
RSVP extension.....	733
rsvp-te statement.....	1008
usage guidelines.....	482
Russian dolls bandwidth model.....	322
S	
Sampling.....	400
sampling.....	936
scope of LSPs.....	29
SE (reservation style).....	466
secondary	
LSPs.....	214, 267, 937
paths.....	45
revert timer.....	215
selection.....	216
secondary path	
verifying.....	780
secondary statement.....	937, 944
usage guidelines.....	214
select statement.....	938
usage guidelines.....	216
send (tracing flag modifier)	
LDP.....	1079
LMP.....	1121
RSVP.....	1013
server (tracing flag).....	1120
session messages, LDP.....	522
session protection, LDP	
configuration.....	631
overview.....	525
session statement.....	1074
usage guidelines.....	630
session-protection statement.....	1075
usage guidelines.....	631
sessions, RSVP.....	472
setup priority, signaled LSPs.....	250
setup protection, RSVP.....	480
setup-protection statement.....	1008
usage guidelines.....	480
shared explicit reservation style.....	466
shared reservations.....	466
shared risk link group	
overview.....	80
show (tracing flag)	
LMP.....	1120
show configuration interfaces command	793
show configuration protocols mpls command	792
show connections command.....	1170, 1356
show ldp database command.....	1310
show ldp fec-filters command.....	1319
show ldp interface command.....	1320
show ldp neighbor command.....	1322
show ldp overview command.....	1324
show ldp path command.....	1329
show ldp route command.....	1331
show ldp session command.....	1335
show ldp statistics command.....	1341
show ldp traffic-statistics command.....	1345
show link-management command.....	1173
show link-management peer command.....	1177
show link-management routing command.....	1179
show link-management statistics command.....	1182
show link-management te-link command.....	1184
show mpls admin-groups command.....	1186
show mpls call-admission-control command.....	1188
show mpls container-lsp command.....	1190
show mpls cspf command.....	1199
show mpls diffserv-te command.....	1201
show mpls egress-protection	
command.....	1197, 1203
show mpls interface command	759, 1205
show mpls label usage command.....	1207
show mpls lsp autobandwidth command.....	1227
show mpls lsp command	761, 805, 1209
show mpls lsp extensive command	
link protection	768
MPLS layer	785
MPLS protocol	795
node-link protection	763

show mpls lsp extensive ingress command		
primary path	780	
show mpls lsp name command.....	804	
show mpls lsp name extensive command.....	804	
show mpls path command.....	1230	
show mpls static-lsp command.....	1233	
show path-computation-client active-pce		
command.....	1378	
show path-computation-client statistics		
command.....	1382	
show performance-monitoring mpls lsp.....	1236	
show route ccc command.....	1359	
show route command	789	
show route forwarding-table command.....	1360	
show route protocol rsvp detail command		
bandwidth load balancing	809	
show route table mpls.0 command	788	
show rsvp interface command.....	1271	
link protection	772	
node-link protection	764	
primary path	780	
show rsvp interface extensive command		
node-link protection	765	
show rsvp neighbor command.....	1276	
show rsvp session command.....	1281	
show rsvp session detail command		
link protection	770	
node-link protection	766	
show rsvp statistics command.....	1291	
show rsvp version command.....	1295	
show security keychain command.....	1349	
show ted database command.....	1242	
show ted link command.....	1250	
show ted protocol command.....	1253	
signal-bandwidth statement.....	938	
usage guidelines.....	688	
signaled LSPs.....	919	
adaptive rerouting.....	249, 825	
administrative groups		
admin-groups statement.....	831	
configuring.....	240	
fast reroute.....	226	
associating addresses.....	228, 880	
configuration statements.....	884	
constrained-path computation		
disabling.....	239, 903, 951	
CoS values.....	244	
damping LSP transitions.....	268	
egress router address.....	212, 953	
fast reroute.....	226, 842, 870	
fate-sharing.....	62, 871	
hold time.....	268, 835	
ingress router address.....	212, 872	
metrics.....	230, 231, 898	
MPLS routers, configuring.....	68	
named paths.....	60	
path		
bandwidth.....	256	
connection retry information.....	229, 932	
length.....	256, 875, 987	
preemption.....	250, 927	
preference levels.....	243, 925	
primary.....	214, 926	
priorities.....	250, 927	
recording routes.....	244	
reoptimization.....	251, 913, 916	
RSVP See RSVP		
secondary.....	214, 937	
standby state.....	267, 944	
tie-breaking rules.....	32, 232, 930	
TTL decrementing.....	236, 904, 907	
signaling component, traffic engineering.....	7	
signaling extensions, RSVP.....	7	
size option to traceoptions statement		
LMP.....	1121	
smart-optimize-timer statement.....	939	
usage guidelines.....	255	
SNMP traps		
MPLS.....	344, 890	
soft-preemption statement		
MPLS.....	940	
usage guidelines.....	238	
RSVP.....	1009	
usage guidelines.....	238	
special labels.....	25	
splitting-merging.....	941	
srlg.....	80	
excluding, common links, secondary		
path.....	90, 95, 116	
overview.....	80	
usage guidelines.....	81	
srlg statement.....	942	
srlg-cost		
usage guidelines.....	81	
srlg-cost statement.....	943	
srlg-value		
usage guidelines.....	81	
srlg-value statement.....	943	

stacked labels.....	26	swap and push (label operation).....	27
standby secondary paths.....	45	swap statement	
standby state, signaled LSPs.....	267, 944	MPLS.....	948
standby statement.....	944	usage guidelines.....	274
usage guidelines.....	267	switch-away-lsps statement.....	949
state (tracing flag)		usage guidelines.....	479
LDP.....	1079	switching-type statement.....	950
LMP.....	1121	usage guidelines.....	688
MPLS.....	955	syntax conventions.....	xxxix
RSVP.....	1013	system log messages	
stateful PCE		MPLS.....	344, 890
active.....	737		
functions.....	737	T	
passive.....	737	targeted hello messages.....	521
stateless PCE.....	737	targeted hello messages, LDP.....	1038
static (tracing flag).....	955	Targeted LDP	
static LSPs		multicast LDP support.....	551
configuring.....	271	targeted-hello statement.....	1077
egress routers.....	274, 276, 882	usage guidelines.....	529, 530
ingress routers.....	271	TCC	
intermediate routers.....	274, 882	configuration.....	659
link protection.....	509	connections.....	663
MPLS, displaying.....	1233	encapsulation.....	659
node protection.....	509	graceful restart	
overview.....	29	configuration.....	665
revert timer.....	215, 275	overview.....	664
static LSP metric.....	231, 898	Layer 2 switching.....	659
static MPLS.....	271	overview.....	644
static routes		te-class-matrix statement.....	952
point-to-multipoint LSPs.....	277	usage guidelines.....	317
static-label-switched path statement		te-link statement.....	1119
usage guidelines.....	271	LMP traffic engineering link	
static-label-switched-path statement		usage guidelines.....	678
static LSP.....	945, 1010	traffic engineering link associated with peer	
statistics		usage guidelines.....	684
MPLS traffic.....	342, 947	usage guidelines.....	725
output file.....	342	technical support	
statistics statement.....	947	contacting JTAC.....	xli
usage guidelines.....	342	TED See traffic engineering database	
strict explicit routes.....	6, 278	export.....	12
strict-targeted-hellos statement.....	1077	import.....	11
usage guidelines.....	531	ted-export (tracing flag).....	955
subscribing to bandwidth.....	1011	ted-import (tracing flag).....	955
subscription statement.....	1011	temporary GMPLS LSP deletion.....	690
usage guidelines.....	323	Throughput	
summary LSA.....	339	definition.....	51
support, technical See technical support		tie-breaking rules, path calculation.....	32, 232, 930
Supported features.....	15, 56, 408, 698	timer (tracing flag)	
swap (label operation).....	27	MPLS.....	955

timer, LDP.....	528, 1038	error	
timers		LDP.....	1078
RSVP.....	483, 1006	MPLS.....	955
Timestamp format.....	52	RSVP.....	1012
to statement		event	
MPLS.....	953	LDP.....	1078
usage guidelines.....	212	RSVP.....	1012
static LSPs		graceful-restart.....	955
usage guidelines (ingress router).....	271	hello-packets	
tracoptions statement		LMP.....	1120
LDP.....	1078	init	
usage guidelines.....	637	LMP.....	1120
LMP.....	1120	initialization.....	1079
usage guidelines.....	687	label.....	1079
MPLS.....	954	lmp.....	1012
usage guidelines.....	361	lsp-history.....	955
RSVP.....	1012	lsping.....	955
usage guidelines.....	491	notification.....	1079
traceroute mpls bgp command.....	1255	nsr-synchronization.....	955
traceroute mpls ldp.....	1352	nsr-synchronization-detail.....	955
traceroute mpls ldp command.....	1352	packets	
traceroute mpls rsvp.....	1298	LDP.....	1079
traceroute mpls rsvp command.....	1298	LMP.....	1120
tracing flag modifiers		RSVP.....	1012
detail		parse	
LDP.....	1079	LMP.....	1120
LMP.....	1121	path	
RSVP.....	1013	LDP.....	1079
disable.....	1079	RSVP.....	1012
receive		pathtear.....	1012
LDP.....	1079	periodic.....	1079
LMP.....	1121	process.....	1120
RSVP.....	1013	resv.....	1013
send		resvtear.....	1013
LDP.....	1079	route.....	1013
LMP.....	1121	LDP.....	1079
RSVP.....	1013	route-socket	
tracing flags		LMP.....	1120
address.....	1078	routing.....	1120
all.....	954	server.....	1120
LMP.....	1120	show	
RSVP.....	1012	LMP.....	1120
automatic bandwidth.....	954	state	
binding.....	1078	LDP.....	1079
connection.....	954	LMP.....	1121
connection-detail.....	954	MPLS.....	955
cspf.....	955	RSVP.....	1013
cspf-link.....	955		
cspf-node.....	955		

- static.....955
 - timer
 - MPLS.....955
 - tracing operations
 - LDP.....637, 1078
 - LMP.....687, 1120
 - MPLS.....361, 954
 - RSVP.....491, 1012
 - track-igp-metric statement.....1080
 - usage guidelines.....628
 - traffic
 - policing.....643
 - protection, MPLS.....45
 - statistics.....342, 947
 - Traffic class.....52
 - traffic engineering.....731
 - BGP destinations.....40
 - fate-sharing.....34
 - IGP and BGP destinations.....41
 - IGP shortcuts.....34
 - information distribution component.....5
 - inter-area, OSPF.....339
 - links.....677
 - LSP metric advertisement.....339
 - overview.....3, 24
 - packet-forwarding component.....4
 - path-selection component.....6
 - routing options.....7
 - signaling component.....7
 - srlg.....80
 - traffic engineering database accuracy.....68
 - Traffic engineering
 - challenges.....383
 - Traffic engineering database
 - export.....12
 - import.....11
 - traffic engineering database.....29
 - accuracy.....68
 - bandwidth update threshold.....476
 - database entries, displaying.....1242
 - link information, displaying.....1250
 - protocols learned from, displaying.....1253
 - traffic-engineered LSPs
 - fast reroute.....327
 - traffic-engineering
 - database
 - export.....850, 866 See credibility
 - import.....876
 - traffic-engineering statement
 - BGP.....963
 - bgp-igp option.....337
 - bgp-igp-both-ribs option.....337
 - MPLS.....962
 - usage guidelines.....336
 - mpls-forwarding option.....338
 - usage guidelines.....285
 - traffic-statistics statement.....1081
 - usage guidelines.....635
 - transit router
 - show route table mpls.0 command788
 - transit routers.....28
 - transit statement
 - static LSP.....1014
 - transit-delay statement.....1122
 - usage guidelines.....686
 - transit-lsp-association
 - usage guidelines.....136
 - transit-lsp-association statement
 - MPLS-TP.....964
 - transitions
 - advertising.....268, 835
 - damping.....268
 - translational cross-connect See TCC
 - transport-address statement.....1083
 - usage guidelines.....539
 - traps, SNMP See SNMP traps
 - TTL decrementing
 - disabling.....236, 904, 907
 - tunnel-services statement.....1015
 - usage guidelines.....490
 - tunneling, MPLS
 - CCC encapsulation.....656
 - example configurations.....658
 - overview.....655
 - RSVP LSPs.....523, 629, 887
 - RSVP LSPs, heterogeneous networks.....630
- ## U
- ultimate-hop popping.....335
 - point-to-multipoint LSPs.....490
 - ultimate-hop-popping statement.....965, 1016
 - usage guidelines.....222, 486
 - uneven load balancing
 - show route protocol rsvp detail
 - command.....809
 - verifying.....808
 - unnumbered interfaces, RSVP.....477

unstable LSPs

fate-sharing See fate-sharing

Unsupported features.....15, 56, 408, 698

update-threshold statement.....1017

usage guidelines.....476

upstream-label statement.....1123

V

verification

BGP session flap prevention.....161

network interfaces.....302

version

RSVP, displaying.....1295

version statement

liveness detection.....1084

VLAN LSP

associated bidirectional packet LSP.....696

configuration overview.....695

configuring.....698

functionality.....694

LSP hierarchy.....695

make-before-break.....697

need.....692

path specification.....695

signaling.....692

supported features.....698

understanding.....692

unsupported features.....698

W

wildcard filter (WF) reservation style.....466

wildcard senders, RSVP.....466

world-readable option to statistics statement

MPLS.....947

world-readable option to traceoptions statement

LDP.....1080

LMP.....1121

MPLS.....955

RSVP.....1013