



---

Junos<sup>®</sup> OS

# MPLS Configuration Guide

Release  
12.1



---

Published: 2012-03-13

Juniper Networks, Inc.  
1194 North Mathilda Avenue  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

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.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

#### *Junos® OS MPLS Configuration Guide*

12.1

Copyright © 2012, 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 . . . . .	xvii
	Documentation and Release Notes . . . . .	xvii
	Using the Examples in This Manual . . . . .	xvii
	Merging a Full Example . . . . .	xviii
	Merging a Snippet . . . . .	xviii
	Documentation Conventions . . . . .	xix
	Documentation Feedback . . . . .	xx
	Requesting Technical Support . . . . .	xxi
	Self-Help Online Tools and Resources . . . . .	xxi
	Opening a Case with JTAC . . . . .	xxii
<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 1</b>	<b>Introduction to Traffic Engineering . . . . .</b>	<b>3</b>
	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 Planning and Analysis . . . . .	6
	Signaling Component . . . . .	7
	Flexible LSP Calculation and Configuration . . . . .	7
<b>Chapter 2</b>	<b>Introduction to MPLS . . . . .</b>	<b>9</b>
	MPLS Introduction . . . . .	10
	MPLS and Traffic Engineering . . . . .	10
	Label Description . . . . .	11
	Special Labels . . . . .	11
	Label Allocation . . . . .	12
	Operations on Labels . . . . .	13
	Routers in an LSP . . . . .	14
	How a Packet Travels Along an LSP . . . . .	14
	Types of LSPs . . . . .	15
	Scope of LSPs . . . . .	15
	Constrained-Path LSP Computation . . . . .	15
	How CSPF Selects a Path . . . . .	17
	Path Selection Tie-Breaking . . . . .	18
	Computing Paths Offline . . . . .	18
	LSPs on an Overloaded Router . . . . .	19

	Fate Sharing .....	19
	SRLG Overview .....	20
	IGP Shortcuts .....	21
	Enabling IGP Shortcuts .....	22
	LSPs Qualified in Shortcut Computations .....	22
	IGP Shortcut Applications .....	23
	IGP Shortcuts and Routing Table .....	23
	IGP Shortcuts and VPNs .....	24
	Advertising LSPs into IGP .....	24
	IP and MPLS Packets on Aggregated Interfaces .....	25
	MPLS Applications .....	26
	BGP Destinations .....	26
	IGP and BGP Destinations .....	28
	Selecting a Forwarding LSP Next Hop .....	28
	MPLS and Routing Tables .....	29
	MPLS and Traffic Protection .....	30
	Fast Reroute Overview .....	31
	Detour Merging Process .....	33
	Detour Computations .....	34
	Fast Reroute Path Optimization .....	34
	Automatic Bandwidth Allocation .....	35
	Point-to-Multipoint LSPs Overview .....	35
<b>Chapter 3</b>	<b>Introduction to DiffServ-Aware Traffic Engineering Configuration Guidelines .....</b>	<b>39</b>
	DiffServ-Aware Traffic Engineering Introduction .....	39
	DiffServ-Aware Traffic Engineering Features .....	40
	DiffServ-Aware Traffic Engineered LSPs .....	40
	DiffServ-Aware Traffic Engineered LSPs Overview .....	40
	DiffServ-Aware Traffic Engineered LSPs Operation .....	41
	Multiclass LSPs .....	42
	Multiclass LSP Overview .....	42
	Establishing a Multiclass LSP on the Differentiated Services Domain .....	43
	Bandwidth Oversubscription Overview .....	43
	LSP Size Oversubscription .....	44
	Link Size Oversubscription .....	44
	Class Type Oversubscription and Local Oversubscription Multipliers .....	45
<b>Part 2</b>	<b>Configuration</b>	
<b>Chapter 4</b>	<b>MPLS Router Configuration Guidelines .....</b>	<b>49</b>
	Configuring the Ingress Router for MPLS-Signaled LSPs .....	49
	Creating Named Paths .....	50
	Examples: Creating Named Paths .....	51
	Configuring Alternate Backup Paths Using Fate Sharing .....	51
	Configuring Fate Sharing .....	52
	Implications for CSPF .....	53

Example: Configuring Fate Sharing . . . . .	53
Example: Configuring a Constrained-Path LSP for Which Junos OS Makes All Forwarding Decisions . . . . .	54
Example: Configuring an Explicit-Path LSP . . . . .	54
Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most Forwarding Decisions and Considers Hop Constraints . . . . .	55
Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most Forwarding Decisions and the Secondary Path Is Explicit . . . . .	56
Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs . . . . .	57
Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages . . . . .	57
PathErr Messages . . . . .	57
Identifying the Problem Link . . . . .	58
Configuring the Router to Improve Traffic Engineering Database Accuracy . . . . .	58
Configuring MPLS-Signaled LSPs to Use GRE Tunnels . . . . .	59
Example: Configuring MPLS-Signaled LSPs to Use GRE Tunnels . . . . .	59
Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks . . . . .	60
Configuring ICMP Message Tunneling for MPLS . . . . .	69
Example: Configuring SRLG . . . . .	69
Example: Excluding SRLG Links Completely for the Secondary LSP . . . . .	78
Example: Configuring SRLG With Link Protection . . . . .	83
Example: Configuring SRLG With Link Protection With the exclude-srlg Option . . . . .	103
Configuring the MPLS Transport Profile for OAM . . . . .	122
MPLS Transport Profile Overview . . . . .	122
Example: Configuring the MPLS Transport Profile for OAM . . . . .	123
<b>Chapter 5 MPLS-Signaled LSP Configuration Guidelines . . . . .</b>	<b>135</b>
LSP Configuration Overview . . . . .	136
Configuring the Ingress and Egress Router Addresses for LSPs . . . . .	139
Configuring the Ingress Router Address for LSPs . . . . .	139
Configuring the Egress Router Address for LSPs . . . . .	139
Preventing the Addition of Egress Router Addresses to Routing Tables . . . . .	140
Configuring Primary and Secondary LSPs . . . . .	141
Configuring Primary and Secondary Paths for an LSP . . . . .	141
Configuring the Revert Timer for LSPs . . . . .	142
Specifying the Conditions for Path Selection . . . . .	143
Configuring a Text Description for LSPs . . . . .	144
Configuring Fast Reroute . . . . .	144
Configuring the Optimization Interval for Fast Reroute Paths . . . . .	146
Adding LSP-Related Routes to the inet.3 Routing Table . . . . .	146
Configuring the Connection Between Ingress and Egress Routers . . . . .	147
Configuring LSP Metrics . . . . .	148
Configuring Dynamic LSP Metrics . . . . .	148
Configuring Static LSP Metrics . . . . .	148
Configuring CSPF Tie Breaking . . . . .	149
Configuring Load Balancing Based on MPLS Labels . . . . .	150
Disabling Normal TTL Decrementing . . . . .	153

Configuring MPLS Soft Preemption . . . . .	154
Configuring Automatic Bandwidth Allocation for LSPs . . . . .	155
Configuring MPLS Statistics for Automatic Bandwidth Allocation . . . . .	156
Configuring Automatic Bandwidth Allocation on LSPs . . . . .	157
Configuring the Automatic Bandwidth Allocation Interval . . . . .	157
Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth . . . . .	158
Configuring the Automatic Bandwidth Adjustment Threshold . . . . .	158
Configuring a Limit on Bandwidth Overflow and Underflow Samples . . . . .	159
Configuring Passive Bandwidth Utilization Monitoring . . . . .	161
Requesting Automatic Bandwidth Allocation Adjustment . . . . .	161
Disabling Constrained-Path LSP Computation . . . . .	162
Configuring Administrative Groups . . . . .	163
Configuring Extended Administrative Groups . . . . .	165
Configuring Preference Values for LSPs . . . . .	167
Disabling Path Route Recording . . . . .	167
Configuring Class of Service for MPLS LSPs . . . . .	167
Class of Service for MPLS Overview . . . . .	168
Configuring the MPLS CoS Bits . . . . .	168
Rewriting IEEE 802.1p Packet Headers with the MPLS CoS Value . . . . .	169
Configuring Adaptive LSPs . . . . .	170
Configuring Priority and Preemption for LSPs . . . . .	171
Optimizing Signaled LSPs . . . . .	172
Configuring the Smart Optimize Timer . . . . .	175
Limiting the Number of Hops in LSPs . . . . .	176
Configuring the Bandwidth Value for LSPs . . . . .	176
Configuring Hot Standby of Secondary Paths . . . . .	177
Damping Advertisement of LSP State Changes . . . . .	178
<b>Chapter 6</b>	
<b>DiffServ-Aware Traffic Engineering Configuration Guidelines . . . . .</b>	<b>179</b>
Configuring the Bandwidth Subscription Percentage for LSPs . . . . .	179
Constraints on Configuring Bandwidth Subscription . . . . .	180
Configuring LSPs for DiffServ-Aware Traffic Engineering . . . . .	181
Configuring Class of Service for the Interfaces . . . . .	181
Configuring IGP . . . . .	182
Configuring Traffic-Engineered LSPs . . . . .	182
Configuring Policing for LSPs . . . . .	183
Configuring Fast Reroute for Traffic-Engineered LSPs . . . . .	183
Configuring Multiclass LSPs . . . . .	184
Configuring Class of Service for the Interfaces . . . . .	184
Configuring the IGP . . . . .	185
Configuring Class-Type Bandwidth Constraints for Multiclass LSPs . . . . .	185
Configuring Policing for Multiclass LSPs . . . . .	186
Configuring Fast Reroute for Multiclass LSPs . . . . .	186

<b>Chapter 7</b>	<b>Static and Explicit-Path LSP Configuration Guidelines . . . . .</b>	<b>189</b>
	Configuring Static LSPs . . . . .	189
	Configuring the Ingress Router for Static LSPs . . . . .	189
	Example: Configuring the Ingress Router . . . . .	191
	Configuring the Intermediate (Transit) and Egress Routers for Static LSPs . . . . .	192
	Example: Configuring an Intermediate Router . . . . .	193
	Example: Configuring an Egress Router . . . . .	194
	Configuring a Bypass LSP for the Static LSP . . . . .	194
	Configuring the Protection Revert Timer for Static LSPs . . . . .	195
	Configuring Static Unicast Routes for Point-to-Multipoint LSPs . . . . .	195
	Configuring Explicit-Path LSPs . . . . .	196
<b>Chapter 8</b>	<b>Point-to-Multipoint LSP Configuration Guidelines . . . . .</b>	<b>199</b>
	Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs . . . . .	199
	Configuring the Primary Point-to-Multipoint LSP . . . . .	200
	Configuring a Branch LSP for Point-to-Multipoint LSPs . . . . .	200
	Configuring the Branch LSP as a Dynamic Path . . . . .	201
	Configuring the Branch LSP as a Static Path . . . . .	201
	Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP . . . . .	201
	Configuring Inter-domain P2MP LSPs . . . . .	219
	Configuring Link Protection for Point-to-Multipoint LSPs . . . . .	221
	Configuring Graceful Restart for Point-to-Multipoint LSPs . . . . .	221
	Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs . . . . .	222
	Example: Configuring Multicast RPF Check Policy for a Point-to-Multipoint LSP . . . . .	223
	Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs . . . . .	223
	Enabling Point-to-Point LSPs to Monitor Egress PE Routers . . . . .	223
	Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases . . . . .	224
<b>Chapter 9</b>	<b>Miscellaneous MPLS Properties Configuration Guidelines . . . . .</b>	<b>225</b>
	Configuring the Maximum Number of MPLS Labels . . . . .	225
	Configuring MPLS to Pop the Label on the Ultimate-Hop Router . . . . .	227
	Configuring Traffic Engineering for LSPs . . . . .	227
	Using LSPs for Both BGP and IGP Traffic Forwarding . . . . .	228
	Using LSPs for Forwarding in Virtual Private Networks . . . . .	228
	Using RSVP and LDP Routes for Forwarding but Not Route Selection . . . . .	229
	Advertising the LSP Metric in Summary LSAs . . . . .	230
	Enabling Interarea Traffic Engineering . . . . .	230
	Enabling Inter-AS Traffic Engineering for LSPs . . . . .	231
	Inter-AS Traffic Engineering Requirements . . . . .	231
	Inter-AS Traffic Engineering Limitations . . . . .	232
	Configuring OSPF Passive TE Mode . . . . .	233
	Configuring MPLS to Gather Statistics . . . . .	233
	Configuring System Log Messages and SNMP Traps for LSPs . . . . .	234
	Configuring MPLS Firewall Filters and Policers . . . . .	236
	Configuring MPLS Firewall Filters . . . . .	236
	Examples: Configuring MPLS Firewall Filters . . . . .	237

	Configuring Policers for LSPs . . . . .	238
	LSP Policer Limitations . . . . .	239
	Example: Configuring an LSP Policer . . . . .	239
	Configuring Automatic Policers . . . . .	240
	Configuring Automatic Policers for LSPs . . . . .	240
	Configuring Automatic Policers for DiffServ-Aware Traffic Engineering LSPs . . . . .	241
	Configuring Automatic Policers for Point-to-Multipoint LSPs . . . . .	242
	Disabling Automatic Policing on an LSP . . . . .	242
	Example: Configuring Automatic Policing for an LSP . . . . .	242
	Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets . . . . .	243
	Configuring MPLS Rewrite Rules . . . . .	243
	Rewriting the EXP Bits of All Three Labels of an Outgoing Packet . . . . .	243
	Rewriting MPLS and IPv4 Packet Headers . . . . .	244
	Configuring BFD for MPLS IPv4 LSPs . . . . .	245
	Configuring BFD for RSVP-Signaled LSPs . . . . .	245
	Configuring a Failure Action for the BFD Session on an RSVP LSP . . . . .	246
	Pinging LSPs . . . . .	247
	Pinging MPLS LSPs . . . . .	247
	Pinging Point-to-Multipoint LSPs . . . . .	248
	Pinging the Endpoint Address of MPLS LSPs . . . . .	248
	Pinging CCC LSPs . . . . .	248
	Pinging Layer 3 VPNs . . . . .	248
	Support for LSP Ping and Traceroute Commands Based on RFC 4379 . . . . .	248
	Tracing MPLS and LSP Packets and Operations . . . . .	249
<b>Part 3</b>	<b>Administration</b>	
<b>Chapter 10</b>	<b>Complete MPLS Applications Configuration Statements . . . . .</b>	<b>253</b>
	[edit logical-systems] Hierarchy Level . . . . .	253
	[edit protocols connections] Hierarchy Level . . . . .	254
	[edit protocols ldp] Hierarchy Level . . . . .	254
	[edit protocols link-management] Hierarchy Level . . . . .	256
	[edit protocols mpls] Hierarchy Level . . . . .	257
	[edit protocols rsvp] Hierarchy Level . . . . .	261
<b>Chapter 11</b>	<b>MPLS Standards and Support . . . . .</b>	<b>265</b>
	Supported MPLS Standards . . . . .	265
	Link-Layer Support . . . . .	267
<b>Chapter 12</b>	<b>MPLS Router Configuration Guidelines Reference . . . . .</b>	<b>269</b>
	Minimum MPLS Configuration . . . . .	269



<b>Chapter 13</b>	<b>DiffServ-Aware Traffic Engineering Configuration Guidelines</b>	
	<b>Reference</b>	<b>271</b>
	DiffServ-Aware Traffic Engineering Standards	271
	DiffServ-Aware Traffic Engineering Terminology	271
	Configuring Routers for DiffServ-Aware Traffic Engineering	272
	Configuring the Bandwidth Model	274
	Configuring Traffic Engineering Classes	274
	Requirements and Limitations for the Traffic Engineering Class	
	Matrix	276
	Configuring Class of Service for DiffServ-Aware Traffic Engineering	276
	Class Type Bandwidth and the LOM	277
	LOM Calculation for the MAM and Extended MAM Bandwidth Models	277
	LOM Calculation for the Russian Dolls Bandwidth Model	277
	Example: LOM Calculation	278
<b>Chapter 14</b>	<b>Summary of MPLS Configuration Statements</b>	<b>281</b>
	adaptive	281
	adjust-interval	282
	adjust-threshold	282
	adjust-threshold-overflow-limit	283
	adjust-threshold-underflow-limit	283
	admin-down	284
	admin-group (for Interfaces)	284
	admin-group (for LSPs)	285
	admin-group-extended	286
	admin-groups-extended	287
	admin-groups-extended-range	288
	admin-groups	289
	admin-group-extended	290
	admin-groups-extended	291
	advertisement-hold-time	292
	allow-fragmentation	292
	always-mark-connection-protection-tlv	293
	associate-backup-pe-groups	293
	associate-lsp	294
	auto-bandwidth	294
	auto-policing	295
	backup-pe-group	296
	bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)	297
	bandwidth (Static LSP)	298
	bandwidth-model	299
	bandwidth-percent	300
	bfd-liveness-detection	301
	class-of-service	302
	description	303
	diffserv-te	304
	disable	305
	encoding-type	305
	exclude (for Administrative Groups)	306

exclude (for Fast Reroute) .....	306
exclude-srlg .....	307
expand-loose-hop .....	308
explicit-null .....	308
failure-action .....	309
fast-reroute .....	310
fate-sharing .....	311
from .....	312
gpip .....	313
hop-limit .....	314
icmp-tunneling .....	315
include-all (for Administrative Groups) .....	315
include-all (for Fast Reroute) .....	316
include-any (for Administrative Groups) .....	316
include-any (for Fast Reroute) .....	317
ingress .....	318
install .....	319
interface .....	320
ipv6-tunneling .....	320
label-switched-path .....	321
ldp-tunneling .....	323
least-fill .....	323
link-protection (Dynamic LSPs) .....	324
link-protection (Static LSPs) .....	325
log-updown .....	326
lsp-attributes .....	327
maximum-bandwidth .....	327
maximum-labels .....	328
metric .....	329
minimum-bandwidth .....	329
monitor-bandwidth .....	330
most-fill .....	330
mpls .....	330
mpls-tp-mode .....	331
mtu-signaling .....	331
next-hop .....	332
no-cspf .....	333
no-decrement-ttl .....	334
no-install-to-address .....	335
no-propagate-ttl .....	336
no-trap .....	337
node-protection (Static LSP) .....	338
oam .....	339
optimize-aggressive .....	340
optimize-timer .....	341
p2mp .....	342
p2mp-lsp-next-hop .....	342
path .....	343
path-mtu .....	344

policing	344
pop	345
preference	346
primary	347
priority	348
protection-revert-time	349
push	350
random	351
record	352
retry-limit	352
retry-timer	353
revert-timer	354
rpf-check-policy	355
rsvp-error-hold-time	356
secondary	357
select	358
signal-bandwidth	358
smart-optimize-timer	359
soft-preemption	360
srlg	360
srlg-cost	361
srlg-value	361
standby	362
static-label-switched-path	363
statistics	365
swap	366
switch-away-lsps	367
switching-type	368
te-class-matrix	369
to	370
traceoptions	371
traffic-engineering	373
transit-lsp-association	374

## Part 4

## Index

Index	377
-------	-----



# List of Figures

<b>Part 1</b>	<b>Overview</b>	
<b>Chapter 2</b>	<b>Introduction to MPLS</b>	<b>9</b>
	Figure 1: Label Encoding	13
	Figure 2: Class-of-Service Bits	13
	Figure 3: CSPF Computation Process	17
	Figure 4: Aggregation Router A Dual-Homed on Core Routers B and C	19
	Figure 5: Typical SPF Tree, Sourced from Router A	21
	Figure 6: Modified SPF Tree, Using LSP A–D as a Shortcut	21
	Figure 7: Modified SPF Tree, Using LSP A–D and LSP A–E as Shortcuts	22
	Figure 8: IGP Shortcuts	23
	Figure 9: IGP Shortcuts in a Bigger Network	23
	Figure 10: SPF Computations with Advertised LSPs	25
	Figure 11: MPLS Application Topology	27
	Figure 12: How BGP Determines How to Reach Next-Hop Addresses	28
	Figure 13: Routing and Forwarding Tables, traffic-engineering bgp	29
	Figure 14: Routing and Forwarding Tables, traffic-engineering bgp-igp	30
	Figure 15: Detours Established for an LSP Using Fast Reroute	32
	Figure 16: Detour After the Link from Router B to Router C Fails	32
	Figure 17: Detours Merging into Other Detours	33
	Figure 18: Point-to-Multipoint LSPs	36
<b>Part 2</b>	<b>Configuration</b>	
<b>Chapter 4</b>	<b>MPLS Router Configuration Guidelines</b>	<b>49</b>
	Figure 19: IPv6 Networks Linked by MPLS IPv4 Tunnels	61
	Figure 20: MPLS-TP OAM Associated Bidirectional LSPs	125
<b>Chapter 5</b>	<b>MPLS-Signaled LSP Configuration Guidelines</b>	<b>135</b>
	Figure 21: least-fill Load Balancing Algorithm Example	173
<b>Chapter 7</b>	<b>Static and Explicit-Path LSP Configuration Guidelines</b>	<b>189</b>
	Figure 22: Static MPLS Configuration	191
<b>Chapter 8</b>	<b>Point-to-Multipoint LSP Configuration Guidelines</b>	<b>199</b>
	Figure 23: RSVP-Signaled Point-to-Multipoint LSP	202



# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xvii</b>
	Table 1: Notice Icons . . . . .	xix
	Table 2: Text and Syntax Conventions . . . . .	xix
<b>Part 2</b>	<b>Configuration</b>	
<b>Chapter 5</b>	<b>MPLS-Signaled LSP Configuration Guidelines</b> . . . . .	<b>135</b>
	Table 3: MPLS LSP Load Balancing Options . . . . .	151
	Table 4: MPLS CoS Values . . . . .	169
<b>Chapter 9</b>	<b>Miscellaneous MPLS Properties Configuration Guidelines</b> . . . . .	<b>225</b>
	Table 5: Sample Scenarios for Using 3, 4, or 5 MPLS Labels . . . . .	226
<b>Part 3</b>	<b>Administration</b>	
<b>Chapter 13</b>	<b>DiffServ-Aware Traffic Engineering Configuration Guidelines</b>	
	<b>Reference</b> . . . . .	<b>271</b>
	Table 6: Default Values for the Traffic Engineering Class Matrix . . . . .	274





# About the Documentation

- Documentation and Release Notes on page xvii
- Using the Examples in This Manual on page xvii
- Documentation Conventions on page xix
- Documentation Feedback on page xx
- Requesting Technical Support on page xxi

## 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>.

## 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.xml;
    }
  }
}
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 [Junos OS CLI User Guide](#).

## Documentation Conventions

Table 1 on page xix 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.

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b> No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces important new terms.</li> <li>Identifies book names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS System Basics Configuration Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>

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

Convention	Description	Examples
<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@# <b>set system domain-name</b> <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; interface names; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> <li>To configure a stub area, include the <b>stub</b> statement at the [edit protocols ospf area area-id] hierarchy level.</li> <li>The console port is labeled <b>CONSOLE</b>.</li> </ul>
< > (angle brackets)	Enclose optional keywords or variables.	<b>stub</b> <default-metric <i>metric</i> >;
(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.	<b>broadcast   multicast</b>  ( <i>string1</i>   <i>string2</i>   <i>string3</i> )
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Enclose a variable for which you can substitute one or more values.	<b>community name members [</b> <i>community-ids</i> <b>]</b>
Indentation and braces ( { } )	Identify 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.	
<b>J-Web GUI Conventions</b>		
<b>Bold text like this</b>	Represents J-Web graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> <li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li> <li>To cancel the configuration, click <b>Cancel</b>.</li> </ul>
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net), or fill out the documentation feedback form at

<https://www.juniper.net/cgi-bin/docbugreport/> . If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release 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 JNASC 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: <https://www.juniper.net/alerts/>
- 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

# Overview

- [Introduction to Traffic Engineering on page 3](#)
- [Introduction to MPLS on page 9](#)
- [Introduction to DiffServ-Aware Traffic Engineering Configuration Guidelines on page 39](#)





## CHAPTER 1

# Introduction to Traffic Engineering

- [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 Planning and Analysis on page 6](#)
- [Signaling Component on page 7](#)
- [Flexible LSP Calculation and Configuration on page 7](#)

## 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 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](#)

## 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.

## 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.

## Offline 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.

## 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.

## CHAPTER 2

# Introduction to MPLS

- [MPLS Introduction on page 10](#)
- [MPLS and Traffic Engineering on page 10](#)
- [Label Description on page 11](#)
- [Special Labels on page 11](#)
- [Label Allocation on page 12](#)
- [Operations on Labels on page 13](#)
- [Routers in an LSP on page 14](#)
- [How a Packet Travels Along an LSP on page 14](#)
- [Types of LSPs on page 15](#)
- [Scope of LSPs on page 15](#)
- [Constrained-Path LSP Computation on page 15](#)
- [How CSPF Selects a Path on page 17](#)
- [Path Selection Tie-Breaking on page 18](#)
- [Computing Paths Offline on page 18](#)
- [LSPs on an Overloaded Router on page 19](#)
- [Fate Sharing on page 19](#)
- [SRLG Overview on page 20](#)
- [IGP Shortcuts on page 21](#)
- [Enabling IGP Shortcuts on page 22](#)
- [LSPs Qualified in Shortcut Computations on page 22](#)
- [IGP Shortcut Applications on page 23](#)
- [IGP Shortcuts and Routing Table on page 23](#)
- [IGP Shortcuts and VPNs on page 24](#)
- [Advertising LSPs into IGP on page 24](#)
- [IP and MPLS Packets on Aggregated Interfaces on page 25](#)
- [MPLS Applications on page 26](#)
- [BGP Destinations on page 26](#)
- [IGP and BGP Destinations on page 28](#)

- [Selecting a Forwarding LSP Next Hop on page 28](#)
- [MPLS and Routing Tables on page 29](#)
- [MPLS and Traffic Protection on page 30](#)
- [Fast Reroute Overview on page 31](#)
- [Detour Merging Process on page 33](#)
- [Detour Computations on page 34](#)
- [Fast Reroute Path Optimization on page 34](#)
- [Automatic Bandwidth Allocation on page 35](#)
- [Point-to-Multipoint LSPs Overview on page 35](#)

## MPLS Introduction

---

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.

## 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:

- [Label Description on page 11](#)
- [Label Allocation on page 12](#)
- [Routers in an LSP on page 14](#)
- [How a Packet Travels Along an LSP on page 14](#)
- [Types of LSPs on page 15](#)
- [Scope of LSPs on page 15](#)
- [Constrained-Path LSP Computation on page 15](#)
- [LSPs on an Overloaded Router on page 19](#)
- [Fate Sharing on page 19](#)
- [IGP Shortcuts on page 21](#)
- [Advertising LSPs into IGP on page 24](#)

---

## Label Description

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.

---

## Special 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 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.

---

## Label Allocation

In the Junos OS, label values are allocated per router. 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 1 on page 13](#) shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.

Figure 1: Label Encoding





Figure 2 on page 13 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 167.

Figure 2: Class-of-Service Bits



## Operations on 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.

---

## 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.
- 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.

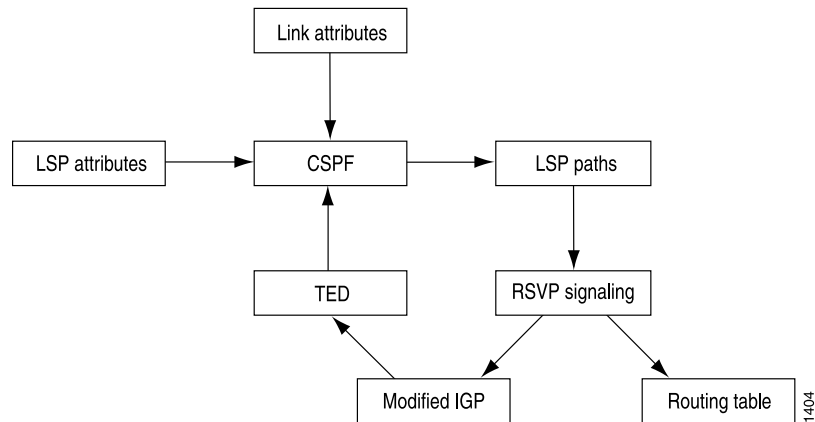
The constraints that CSPF considers include:

- 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 3 on page 17](#) 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 3: CSPF Computation Process



This section discusses the following topics:

- [How CSPF Selects a Path on page 17](#)
- [Path Selection Tie-Breaking on page 18](#)
- [Computing Paths Offline on page 18](#)

## How CSPF Selects a Path

To select a path, CSPF follows these steps:

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 starts with those that have the highest bandwidth requirement.
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. 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's are computed, one from the ingress router to Router A, the other from Router A to the egress router.
6. If several paths have equal cost, chooses the one whose last-hop address is the same as the LSP's destination.
7. If several equal-cost paths remain, selects the one with the fewest number of hops.
8. If several equal-cost paths remain, applies the CSPF load-balancing rule configured on the LSP (least fill, most fill, or random).

## Path Selection Tie-Breaking

---

If more than one path is available after the rules from the previous section have been applied, a tie-breaking rule is applied to choose the path for the LSP. 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.
- 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 rule used depends on the configuration. Random is the default rule.

For the other rules, the following definitions are needed:

- 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

## Computing 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



## 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.

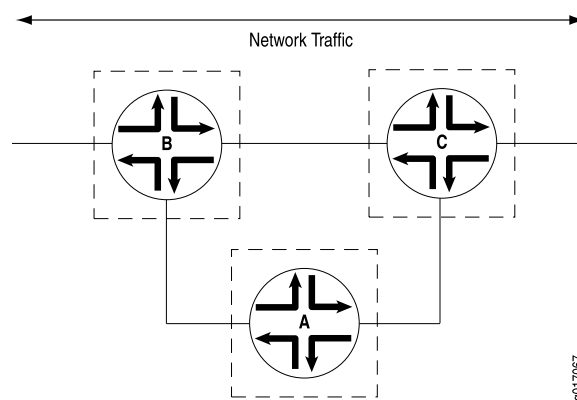
You cannot establish transit LSPs through an overloaded router. However, you can configure ingress and egress LSPs through an overloaded router.



**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 4 on page 19](#), 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 4: Aggregation Router A Dual-Homed on Core Routers B and C**



## 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. For more information about fate sharing, see the [Junos OS Routing Protocols Configuration Guide](#).

## 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 69](#)
- [Example: Excluding SRLG Links Completely for the Secondary LSP on page 78](#)
- [Example: Configuring SRLG With Link Protection on page 83](#)


- [Example: Configuring SRLG With Link Protection With the exclude-srlg Option on page 103](#)
- [Fate Sharing on page 19](#)

## IGP Shortcuts

Link-state protocols, such as OSPF and IS-IS, use the 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. 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 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 5 on page 21](#)).

**Figure 5: 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 6 on page 21](#).

**Figure 6: 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 5 on page 21](#)). If the LSP metric is lower, LSP A–D is used ([Figure 6 on page 21](#)). If both metrics are equal, LSP A–D is chosen because LSP paths 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 7 on page 22](#).

**Figure 7: Modified SPF Tree, Using LSP A–D and LSP A–E as Shortcuts**



This section discusses the following topics:

- [Enabling IGP Shortcuts on page 22](#)
- [LSPs Qualified in Shortcut Computations on page 22](#)
- [IGP Shortcut Applications on page 23](#)
- [IGP Shortcuts and Routing Table on page 23](#)
- [IGP Shortcuts and VPNs on page 24](#)

#### Related Documentation

- [Configuring IS-IS Traffic Engineering Attributes](#)
- [OSPF Extensions to Support 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. For information about enabling IGP shortcuts for IS-IS and OSPF, see the [Junos OS Routing Protocols Configuration Guide](#). 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.

## LSPs Qualified in 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 7 on page 22](#), 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 8 on page 23](#), 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 n2 LSPs from all routers in Region 1 to all routers in Region 2 and allows efficient resource controls on the backbone network.

**Figure 8: IGP Shortcuts**



Shortcuts allow you to deploy LSPs into a network in an incremental, hierarchical fashion. In [Figure 9 on page 23](#), 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 9: 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 Table

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 28](#)), 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 227](#)
- [Configuring IS-IS Traffic Engineering Attributes](#)
- [OSPF Extensions to Support Traffic Engineering](#)


---

## Advertising LSPs into IGP

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 10 on page 25](#).

Figure 10: 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 10 on page 25](#), 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 the [Junos OS Network Interfaces Configuration Guide](#).

---

## 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 26](#)
- [IGP and BGP Destinations on page 28](#)
- [Selecting a Forwarding LSP Next Hop on page 28](#)

---

## BGP Destinations

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

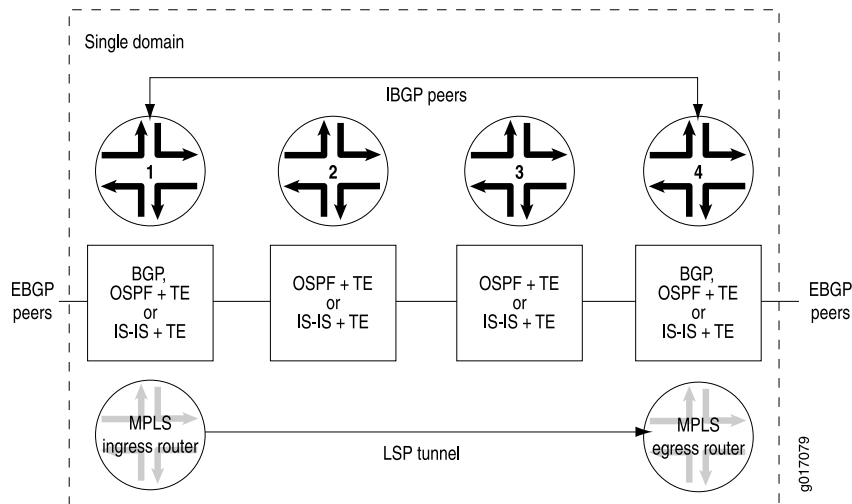
Both IBGP and EBGP take advantage of the LSP host routes without requiring extra configuration. BGP compares the BGP next-hop address with the LSP host route. If a match is found, the packets for the BGP route are label-switched over the LSP. If multiple BGP routes share the same next-hop address, all 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.

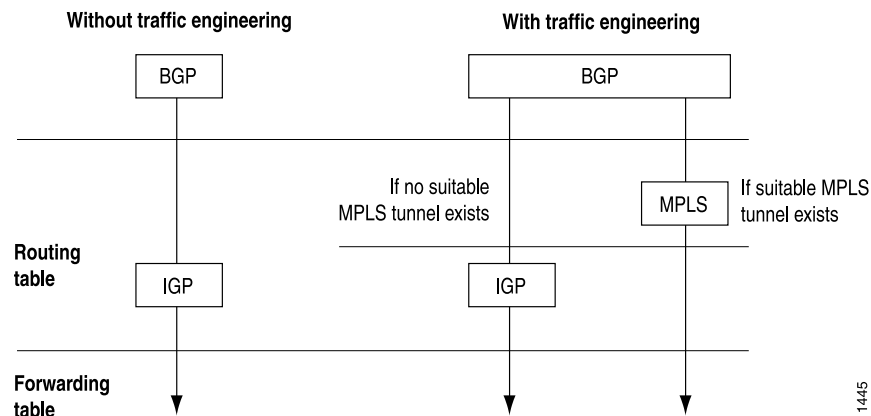
Figure 11 on page 27 shows an MPLS topology that illustrates how MPLS and LSPs work. This topology consists of a single domain with four routers. The two routers at the edges of the domain, Router 1 and Router 4, are running EBGP to communicate with peers outside the domain and IBGP to communicate between themselves. For intradomain communication, all four routers are running an IGP. Finally, an LSP tunnel exists from Router 1 to Router 4.

Figure 11: 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 12 on page 28.) 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 12 on page 28.)

Figure 12: 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 13 on page 29](#)), 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 11 on page 27](#) (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 14 on page 30](#)) 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

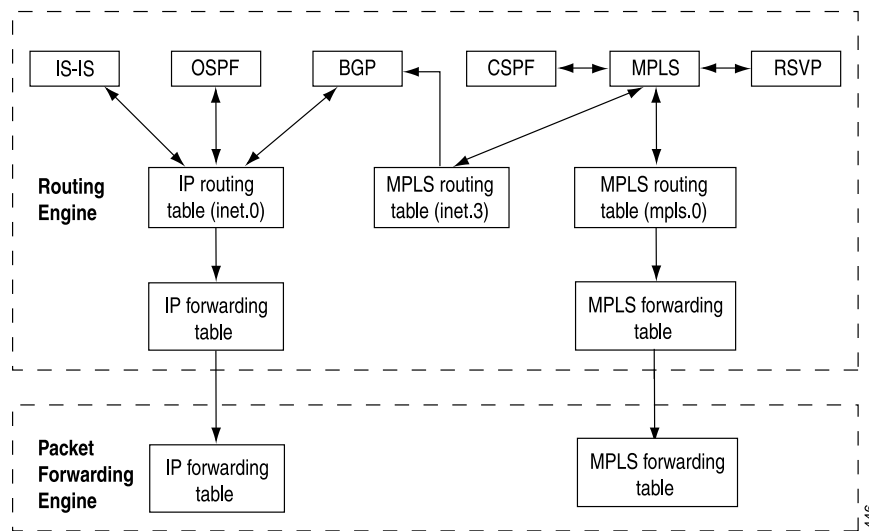
If more than one LSP tunnel to a BGP next hop exists, the prefixes learned from the BGP next hop are randomly divided among the LSP tunnels. To control which LSP BGP uses to forward data for a given prefix, use the **install-nexthop** statement in the export policy

applied to the forwarding table. For more information, see the [Junos OS Routing Protocols Configuration Guide](#).

## 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 13 on page 29 and Figure 14 on page 30 illustrate the routing tables in the two traffic engineering configurations.)

Figure 13: 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.

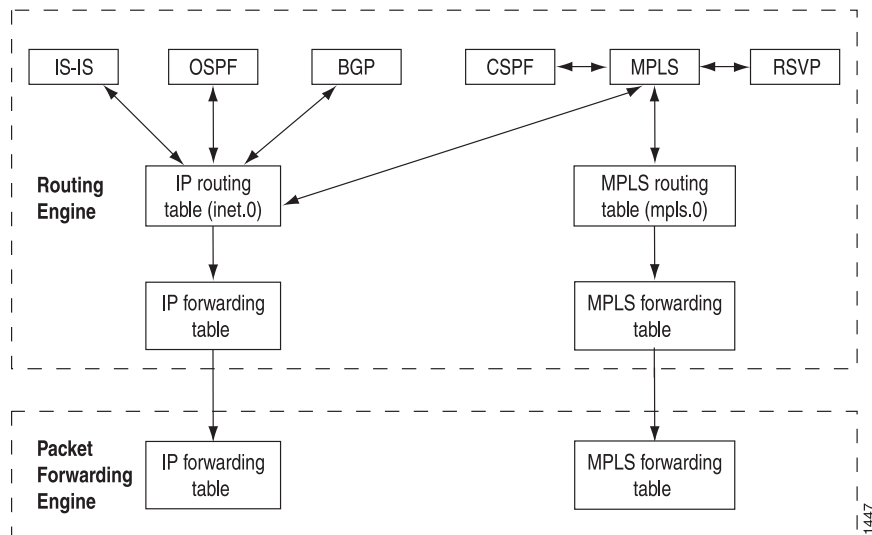
MPLS also maintains an MPLS path routing table (**mpls.0**), which contains a list of the next label-switched router in each LSP. This routing table is used on transit routers to route packets to the next router along an LSP.

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 14: 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” on page 177](#).

- **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 31](#). For information about configuring fast reroute, see [“Configuring Fast Reroute” on page 144](#).
- **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](#).

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.

## 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 15 on page 32](#) 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 that are not shown in the figure.

Fast reroute protects traffic against any single point of failure between the ingress and egress routers. 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 15: 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 16 on page 32](#) illustrates the detour taken when the link between Router B and Router C fails.

**Figure 16: 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 15 on page 32](#) 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 17 on page 33](#), 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 17: Detours Merging into Other Detours**



## 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.

---

## Automatic Bandwidth Allocation

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.

---

## 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 18 on page 36](#). 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 *draft-ietf-mpls-ldp-p2mp-10.txt*, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*.

Figure 18: 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 Configuration Guide](#)
- 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 LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs



## CHAPTER 3

# Introduction to DiffServ-Aware Traffic Engineering Configuration Guidelines

- [DiffServ-Aware Traffic Engineering Introduction on page 39](#)
- [DiffServ-Aware Traffic Engineering Features on page 40](#)
- [DiffServ-Aware Traffic Engineered LSPs on page 40](#)
- [DiffServ-Aware Traffic Engineered LSPs Overview on page 40](#)
- [DiffServ-Aware Traffic Engineered LSPs Operation on page 41](#)
- [Multiclass LSPs on page 42](#)
- [Multiclass LSP Overview on page 42](#)
- [Establishing a Multiclass LSP on the Differentiated Services Domain on page 43](#)
- [Bandwidth Oversubscription Overview on page 43](#)
- [LSP Size Oversubscription on page 44](#)
- [Link Size Oversubscription on page 44](#)
- [Class Type Oversubscription and Local Oversubscription Multipliers on page 45](#)

## 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 238](#).

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 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 40](#)
- [DiffServ-Aware Traffic Engineered LSPs Operation on page 41](#)

## 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 [Junos OS Class of Service Configuration Guide](#).
- For EXP bits, see “Label Allocation” on page 12.
- 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 42](#)
- [Establishing a Multiclass LSP on the Differentiated Services Domain on page 43](#)

## 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 [“Label Allocation” on page 12](#).



For more information about forwarding classes, see the *Junos OS Class of Service Configuration Guide*.

## 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 IGP advertises 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.

## 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 44](#)
- [Link Size Oversubscription on page 44](#)
- [Class Type Oversubscription and Local Oversubscription Multipliers on page 45](#)

---

## 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.

---

## 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 179](#).

## 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 277](#)
  - [LOM Calculation for the MAM and Extended MAM Bandwidth Models on page 277](#)
  - [LOM Calculation for the Russian Dolls Bandwidth Model on page 277](#)
  - [Example: LOM Calculation on page 278](#)
-



## PART 2

# Configuration

- [MPLS Router Configuration Guidelines on page 49](#)
- [MPLS-Signaled LSP Configuration Guidelines on page 135](#)
- [DiffServ-Aware Traffic Engineering Configuration Guidelines on page 179](#)
- [Static and Explicit-Path LSP Configuration Guidelines on page 189](#)
- [Point-to-Multipoint LSP Configuration Guidelines on page 199](#)
- [Miscellaneous MPLS Properties Configuration Guidelines on page 225](#)



## CHAPTER 4

# MPLS Router Configuration Guidelines

- [Configuring the Ingress Router for MPLS-Signaled LSPs on page 49](#)
- [Example: Configuring a Constrained-Path LSP for Which Junos OS Makes All Forwarding Decisions on page 54](#)
- [Example: Configuring an Explicit-Path LSP on page 54](#)
- [Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most Forwarding Decisions and Considers Hop Constraints on page 55](#)
- [Example: Configuring a Constrained-Path LSP for Which Junos OS Makes Most Forwarding Decisions and the Secondary Path Is Explicit on page 56](#)
- [Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs on page 57](#)
- [Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages on page 57](#)
- [Configuring MPLS-Signaled LSPs to Use GRE Tunnels on page 59](#)
- [Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks on page 60](#)
- [Configuring ICMP Message Tunneling for MPLS on page 69](#)
- [Example: Configuring SRLG on page 69](#)
- [Example: Excluding SRLG Links Completely for the Secondary LSP on page 78](#)
- [Example: Configuring SRLG With Link Protection on page 83](#)
- [Example: Configuring SRLG With Link Protection With the exclude-srlg Option on page 103](#)
- [Configuring the MPLS Transport Profile for OAM on page 122](#)

## 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 50](#)
- [Configuring Alternate Backup Paths Using Fate Sharing on page 51](#)

## 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. For more information, see the [Junos OS Routing Protocols and Policies Command Reference](#).

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 52](#)
- [Implications for CSPF on page 53](#)
- [Example: Configuring Fate Sharing on page 53](#)

---

### 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.

---

### 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 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 **10.1.1.32** 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 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 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 269](#) and Minimum RSVP Configuration.

## 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 15](#).) 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 RSVP Interfaces](#).

This section discusses the following topics:

- [PathErr Messages on page 57](#)
- [Identifying the Problem Link on page 58](#)
- [Configuring the Router to Improve Traffic Engineering Database Accuracy on page 58](#)

### 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 Configuration Guide](#).

### 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 60](#)
- [Overview on page 60](#)
- [Configuration on page 63](#)
- [Verification on page 68](#)

### 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 19 on page 61](#), 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 19: 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 19 on page 61](#) 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 [Junos OS CLI User Guide](#)*.

To configure Device PE1:

1. Configure the interfaces.
 

```

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

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

user@PE1# set interfaces 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 63](#).

## 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 49](#)
- [Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs on page 57](#)
- [Minimum RSVP Configuration](#)

## Configuring ICMP Message Tunneling for MPLS

Internet Control Message Protocol (ICMP) is one of the core TCP/IP protocols. Routers can send ICMP advertisements over the network to enable hosts to discover the addresses of operating routers. ICMP is also useful for network debugging by enabling the ping and traceroute functions commonly used by network administrators.

ICMP message tunneling enables you to send ICMP messages over an LSP for debugging purposes. To enable this feature, configure the **icmp-tunneling** statement at the **[edit protocols mpls]** hierarchy level on each of the routers from which you wish to receive ICMP messages.

ICMP messages generated by an intermediate LSR (for an LSP) are forwarded to the egress router using the same LSP. The egress router uses IP routing or a reverse LSP to send the ICMP message back towards the original source of the packet that generated the ICMP message at the intermediate LSR.

ICMP message tunneling can handle ICMP traceroute using ICMP time exceeded messages. It can also handle path MTU discovery which relies on the ICMP error message “fragmentation needed, but do-not-fragment bit is set.”

To configure ICMP message tunneling for MPLS, include the **icmp-tunneling** statement:

```
icmp-tunneling;
```

You can include this statement at the following hierarchy levels:

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

## Example: Configuring SRLG

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

- [Requirements on page 70](#)
- [Overview on page 70](#)
- [Configuration on page 71](#)
- [Verification on page 76](#)

## 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 [Junos OS 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.

```
[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 76](#)
- [Verify TE-Link SRLG on page 76](#)
- [Verify Standby Secondary Path on page 77](#)

### 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 20](#)
  - [Example: Excluding SRLG Links Completely for the Secondary LSP on page 78](#)
  - [srlg on page 360](#)
  - [srlg-cost on page 361](#)
  - [srlg-value on page 361](#)

---

## 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 78](#)
- [Overview on page 78](#)
- [Configuration on page 79](#)
- [Verification on page 82](#)

### 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 [Junos OS 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 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 20](#)
- [Example: Configuring SRLG on page 69](#)
- [exclude-srlg on page 307](#)

## Example: Configuring SRLG With Link Protection

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

- [Requirements on page 84](#)
- [Overview on page 84](#)
- [Configuration on page 84](#)
- [Verification on page 101](#)

## 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 [Junos OS 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-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 [Junos OS 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 **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 [Junos OS 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 **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 [Junos OS 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 [Junos OS 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 [Junos OS 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 [Junos OS 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 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
  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) 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 20](#)
- [Example: Configuring SRLG on page 69](#)
- [Example: Configuring SRLG With Link Protection With the exclude-srlg Option on page 103](#)

## 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 103](#)
- [Overview on page 103](#)
- [Configuration on page 104](#)
- [Verification on page 121](#)

### 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 [Junos OS 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 [Junos OS 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 [Junos OS 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 *Junos OS 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 [Junos OS 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 [Junos OS 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 [Junos OS 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 20](#)
  - [Example: Configuring SRLG on page 69](#)
  - [Example: Configuring SRLG With Link Protection on page 83](#)
  - [exclude-srlg on page 307](#)

---

## Configuring the MPLS Transport Profile for OAM

- [MPLS Transport Profile Overview on page 122](#)
- [Example: Configuring the MPLS Transport Profile for OAM on page 123](#)

### 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*.

### 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 123](#)
- [Overview on page 123](#)
- [Configuration on page 125](#)
- [Verification on page 132](#)

---

#### 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 20 on page 125 shows the topology used in this example.

Figure 20: MPLS-TP OAM Associated Bidirectional LSPs



g0-40928

## 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
Explct 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.

## CHAPTER 5

# MPLS-Signaled LSP Configuration Guidelines

- [LSP Configuration Overview on page 136](#)
- [Configuring the Ingress and Egress Router Addresses for LSPs on page 139](#)
- [Configuring Primary and Secondary LSPs on page 141](#)
- [Configuring a Text Description for LSPs on page 144](#)
- [Configuring Fast Reroute on page 144](#)
- [Configuring the Optimization Interval for Fast Reroute Paths on page 146](#)
- [Adding LSP-Related Routes to the inet.3 Routing Table on page 146](#)
- [Configuring the Connection Between Ingress and Egress Routers on page 147](#)
- [Configuring LSP Metrics on page 148](#)
- [Configuring CSPF Tie Breaking on page 149](#)
- [Configuring Load Balancing Based on MPLS Labels on page 150](#)
- [Disabling Normal TTL Decrementing on page 153](#)
- [Configuring MPLS Soft Preemption on page 154](#)
- [Configuring Automatic Bandwidth Allocation for LSPs on page 155](#)
- [Disabling Constrained-Path LSP Computation on page 162](#)
- [Configuring Administrative Groups on page 163](#)
- [Configuring Extended Administrative Groups on page 165](#)
- [Configuring Preference Values for LSPs on page 167](#)
- [Disabling Path Route Recording on page 167](#)
- [Configuring Class of Service for MPLS LSPs on page 167](#)
- [Configuring Adaptive LSPs on page 170](#)
- [Configuring Priority and Preemption for LSPs on page 171](#)
- [Optimizing Signaled LSPs on page 172](#)
- [Configuring the Smart Optimize Timer on page 175](#)
- [Limiting the Number of Hops in LSPs on page 176](#)
- [Configuring the Bandwidth Value for LSPs on page 176](#)

- [Configuring Hot Standby of Secondary Paths on page 177](#)
- [Damping Advertisement of LSP State Changes on page 178](#)

---

## LSP Configuration Overview

To configure an MPLS-signaled LSP, you define the properties associated with the LSP on the ingress router. Include the **label-switched-path** statement:

```
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 percent;
    adjust-threshold-overflow-limit number;
    adjust-threshold-underflow-limit number;
    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;
  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 ] | noinclude-any);
  }
  from address;
  hop-limit number;
  install {
    destination-prefix/prefix-length <active>;
  }
  ldp-tunneling;
  link-protection;
  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);
  }
}
```

```

metric number;
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 {
  cleanup-timer seconds;
}
standby;
to address;
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 mpls]
- [edit logical-systems *logical-system-name* protocols mpls]

Each LSP must have a name, *lsp-name*, which can be up to 64 characters long and can contain letters, digits, periods (.), and hyphens (-). The name must be unique within the ingress router. For ease of management and identification, configure unique names across the entire domain.

When you configure LSPs, you can specify the following statements either for each LSP or for each path. For statements that you configure on a per-LSP basis, the value applies to all paths in the LSP. For statements that you configure on a per-path basis, the path value overrides the per-LSP value.

- adaptive
- admin-group (for LSPs)
- auto-bandwidth
- bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)
- class-of-service
- hop-limit
- no-cspf
- optimize-timer
- preference
- priority
- record or no-record
- standby



For maintenance purposes, you can also configure the following attributes across all LSPs and any paths within those LSPs:

- [admin-group](#) (for LSPs)
- [bandwidth](#) (Fast Reroute, Signaled, and Multiclass LSPs)
- [class-of-service](#)
- [no-decrement-ttl](#)
- [no-record](#)
- [optimize-timer](#)
- [preference](#)
- [priority](#)
- [smart-optimize-timer](#)
- [standby](#)

## 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 139](#)
- [Configuring the Egress Router Address for LSPs on page 139](#)
- [Preventing the Addition of Egress Router Addresses to Routing Tables on page 140](#)

### 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 50. 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 141](#)
- [Configuring the Revert Timer for LSPs on page 142](#)
- [Specifying the Conditions for Path Selection on page 143](#)

### 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 147.](#))

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 which 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.

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 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.

To provide a textual description for the LSP, include the **description** statement:

```
description text;
```

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 description text can be no more than 80 characters in length.

---

## 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:

```
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. Once fast reroute is enabled, the ingress router signals all the downstream routers 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](#).

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 182](#).

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” on page 163](#).

## 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 Routing Table

---

By default, a host route toward the egress router is installed in the **inet.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** routing table.

Unlike the routes in the **inet.0** table, routes in the **inet.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** is to permit BGP to perform next-hop resolution. To examine the **inet.3** table, use the **show route table inet.3** command.

To inject additional routes into the **inet.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** 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** routing table.

For BGP next-hop resolution, it makes no difference whether a route is in **inet.0** or **inet.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 148](#)
- [Configuring Static LSP Metrics on page 148](#)

### 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 [“IGP Shortcuts” on page 21](#)), 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 Configuration Guide](#).

## 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 17](#).

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:

- To configure a random tie-breaking rule for CSPF to use to choose among equal-cost paths, include 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*]

---

## Configuring Load Balancing Based on MPLS Labels

Juniper Networks routers can load-balance on a per-packet basis in MPLS. Load balancing can be performed on information in both the IP header and on up to three MPLS labels, providing a more uniform distribution of MPLS traffic to next hops. This feature is enabled on supported platforms by default and requires no 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.

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 in any way, 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).

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.

To load-balance based on the MPLS label information, configure the **family mpls** statement:

```
[edit forwarding-options hash-key]
family mpls {
  label-1;
  label-2;
  label-3;
  no-labels;
  no-label-1-exp;
  payload {
    ether-pseudowire;
```

```

ip {
  layer-3-only;
  port-data {
    destination-lsb;
    destination-msb;
    source-lsb;
    source-msb;
  }
}

```

You can include this statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* forwarding-options hash-key]
- [edit forwarding-options hash-key]

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.

[Table 3 on page 151](#) provides detailed information about all of the possible MPLS LSP load-balancing options.

**Table 3: MPLS LSP Load Balancing Options**

Statement	MPLS LSP Load Balancing Options
<b>label-1</b>	Include the first label in the hash key. Use this option for single label packets.
<b>label-2</b>	Include the second label in the hash key. You must also configure the <b>label-1</b> option. The entire first label and the first 16 bits of the second label are used in the hash key.
<b>label-3</b>	Include the third label in the hash key. You must also configure the <b>label-1</b> option and the <b>label-2</b> option.
<b>no-labels</b>	Excludes MPLS labels from the hash key.
<b>no-label-1-exp</b>	Excludes the EXP bit of the top label from the hash key. You must also configure the <b>label-1</b> 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.
<b>payload</b>	Allows you to configure which parts of the IP packet payload to include in the hash key.
<b>ether-pseudowire</b>	(M120, M320, MX Series, and T Series routers only5)—Load-balance IPv4 traffic over Layer 2 Ethernet pseudowires.
<b>ip</b>	Include the IPv4 or IPv6 address in the hash key. You must also configure either <b>label-1</b> or <b>no-labels</b> .

Table 3: MPLS LSP Load Balancing Options (*continued*)

Statement	MPLS LSP Load Balancing Options
<b>layer-3-only</b>	Include only the Layer 3 IP information in the hash key. Excludes all of the <b>port-data</b> bytes from the hash key.
<b>port-data</b>	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 <b>source-msb</b> , <b>source-lsb</b> , <b>destination-msb</b> , and <b>destination-lsb</b> options at the <code>[edit forwarding-options hash-key family mpls payload ip port-data]</code> hierarchy level. To prevent all four bytes from being hashed, include the <b>layer-3-only</b> statement at the <code>[edit forwarding-options hash-key family mpls payload ip]</code> hierarchy level.
<b>destination-lsb</b>	Include the least significant byte of the destination port in the hash key. Can be combined with any of the other <b>port-data</b> options.
<b>destination-msb</b>	Include the most significant byte of the destination port in the hash key. Can be combined with any of the other <b>port-data</b> options.
<b>source-lsb</b>	Include the least significant byte of the source port in the hash key. Can be combined with any of the other <b>port-data</b> options.
<b>source-msb</b>	Include the most significant byte of the source port in the hash key. Can be combined with any of the other <b>port-data</b> 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, configure the **label-1** 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-1;
payload {
  ip;
}
```

- 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 Core 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 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 Configuration Guide 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 Internet draft draft-ietf-mpls-soft-preemption-02.txt, *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.

## 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.



**NOTE:** You might not be able to use this feature to adjust the bandwidth of fast-reroute LSPs. 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.

To configure automatic bandwidth allocation, complete the steps in the following sections:

- [Configuring MPLS Statistics for Automatic Bandwidth Allocation on page 156](#)
- [Configuring Automatic Bandwidth Allocation on LSPs on page 157](#)
- [Requesting Automatic Bandwidth Allocation Adjustment on page 161](#)

## Configuring MPLS Statistics for Automatic Bandwidth Allocation

To enable automatic bandwidth allocation, you first need to configure MPLS statistics. Include the **auto-bandwidth** option for the **statistics** statement. You can also use the **interval** option to specify the interval for calculating the average bandwidth usage.

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. You can also set the adjustment interval on specific LSPs.



**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 the Automatic Bandwidth Allocation Interval” on page 157](#).

To configure the MPLS and automatic bandwidth allocation statistics, include the **statistics** statement:

```
statistics {  
  auto-bandwidth;  
  file filename <files number> <size size> <world-readable | no-world-readable>;  
  interval seconds;  
}
```

You can include this statement at the following hierarchy levels:

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

## 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;
  minimum-bandwidth bps;
  maximum-bandwidth bps;
  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*]**

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 157](#)
- [Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 158](#)
- [Configuring the Automatic Bandwidth Adjustment Threshold on page 158](#)
- [Configuring a Limit on Bandwidth Overflow and Underflow Samples on page 159](#)
- [Configuring Passive Bandwidth Utilization Monitoring on page 161](#)

### 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 MPLS Statistics for Automatic Bandwidth Allocation” on page 156](#).

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, 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 161](#)).
- 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 161](#)), 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**

- [request mpls lsp adjust-autobandwidth](#)

---

## 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 Configuration Guide](#).
- 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 Configuration Guide](#).

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 15](#).

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 172](#).

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 147](#).

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

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 171](#).

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 17.



**NOTE:** Changing the LSP's administrative group causes an immediate recomputation of the route; therefore, the LSP might be rerouted.

## Configuring Extended Administrative Groups

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 on page 163](#)

---

## 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

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 168](#)
- [Configuring the MPLS CoS Bits on page 168](#)
- [Rewriting IEEE 802.1p Packet Headers with the MPLS CoS Value on page 169](#)

## 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 [“Label Allocation” on page 12](#).

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 [Junos OS Class of Service Configuration Guide](#).

## 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 [Junos OS Class of Service Configuration Guide](#) 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 [Junos OS Class of Service Configuration Guide](#).



**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 169 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 [Junos OS Class of Service Configuration Guide](#).

**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 [Junos OS Class of Service Configuration Guide](#).

## 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 control carefully 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).

Configuring 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 162](#).

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:


- 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.
- For each current (old) path, the software stores the four smallest values for bandwidth availability for the links traversed in ascending order.
- The software also stores the four smallest bandwidth availability values for the new path, corresponding to the links traversed in ascending order.
- 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.
- 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:

- If the new path has a lower IGP metric, it is accepted.
- If the new path has an equal IGP metric and lower hop count, it is accepted.
- If you choose **least-fill** as a load balancing algorithm, LSPs are load balanced as follows:
  - 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%.
  - For **random** or **most-fill** algorithms, this rule does not apply.

The following example illustrates how the **least-fill** load balancing algorithm works.

**Figure 21: least-fill Load Balancing Algorithm Example**



As shown in [Figure 21 on page 173](#), 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 1GE 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.

## Configuring the Smart Optimize Timer

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 30](#).

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 30](#)
- [Optimizing Signaled LSPs on page 172](#)

---

## 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.

## Configuring Hot Standby of Secondary Paths

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 6

# DiffServ-Aware Traffic Engineering Configuration Guidelines

- [Configuring the Bandwidth Subscription Percentage for LSPs on page 179](#)
- [Configuring LSPs for DiffServ-Aware Traffic Engineering on page 181](#)
- [Configuring Multiclass LSPs on page 184](#)

### 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 45](#) 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 [Junos OS Routing Protocols and Policies Command Reference](#).

## 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 272](#)) 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 274](#).

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 181](#)
- [Configuring IGP on page 182](#)
- [Configuring Traffic-Engineered LSPs on page 182](#)
- [Configuring Policing for LSPs on page 183](#)
- [Configuring Fast Reroute for Traffic-Engineered LSPs on page 183](#)

### 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 [Junos OS Class of Service Configuration Guide](#).

## 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 Configuration Guide*.

## 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-proto-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 238](#).

## 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 144](#).

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 272](#)) 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 274](#).

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 184](#)
- [Configuring the IGP on page 185](#)
- [Configuring Class-Type Bandwidth Constraints for Multiclass LSPs on page 185](#)
- [Configuring Policing for Multiclass LSPs on page 186](#)
- [Configuring Fast Reroute for Multiclass LSPs on page 186](#)

### 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 [Junos OS Class of Service Configuration Guide](#).

## 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 Configuration Guide*.

## 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 238](#) and [“Configuring Automatic Policers” on page 240](#).

## 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 182](#).

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 7

# Static and Explicit-Path LSP Configuration Guidelines

- [Configuring Static LSPs on page 189](#)
- [Configuring Explicit-Path LSPs on page 196](#)

## 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 189](#)
- [Configuring the Intermediate \(Transit\) and Egress Routers for Static LSPs on page 192](#)
- [Configuring a Bypass LSP for the Static LSP on page 194](#)
- [Configuring the Protection Revert Timer for Static LSPs on page 195](#)
- [Configuring Static Unicast Routes for Point-to-Multipoint LSPs on page 195](#)

## 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](#))

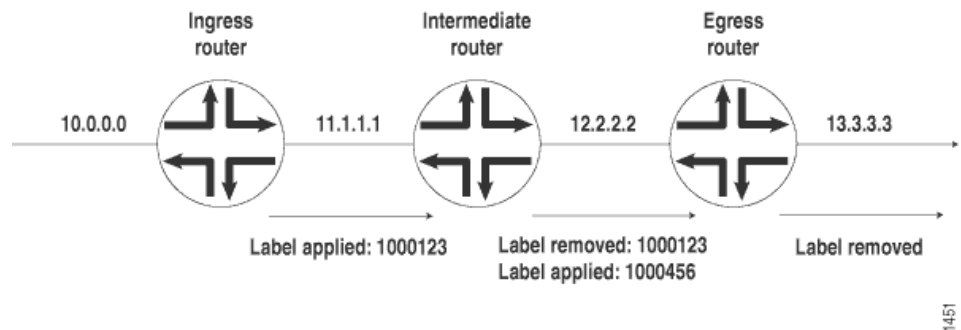
To determine whether a static ingress route is installed, use the command **show route table inet.0 protocol static**. Sample output follows. The **push** keyword denotes that a label is to be added in front of an IP packet.

```
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 22 on page 191](#)).

Figure 22: 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 path1;
    }
  }
}

```

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 142](#).

### 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 35](#), [“Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs” on page 199](#), and [Configuring CCC Switching for Point-to-Multipoint LSPs](#).

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 Configuration Guide](#).

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 162](#), 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 50](#). 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 141](#).
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 162](#).
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 8

# Point-to-Multipoint LSP Configuration Guidelines

- [Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs on page 199](#)
- [Example: Configuring an RSVP-Signaled Point-to-Multipoint LSP on page 201](#)
- [Configuring Inter-domain P2MP LSPs on page 219](#)
- [Configuring Link Protection for Point-to-Multipoint LSPs on page 221](#)
- [Configuring Graceful Restart for Point-to-Multipoint LSPs on page 221](#)
- [Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs on page 222](#)
- [Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs on page 223](#)
- [Enabling Point-to-Point LSPs to Monitor Egress PE Routers on page 223](#)
- [Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases on page 224](#)

### 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 35](#).

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 200](#)
- [Configuring a Branch LSP for Point-to-Multipoint LSPs on page 200](#)

## 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 49](#) 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 172](#) 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 200](#) 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 resigned. 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 201](#)
- [Configuring the Branch LSP as a Static Path on page 201](#)

### 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 189](#) for more information.

## Example: Configuring 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 201](#)
- [Overview on page 201](#)
- [Configuration on page 202](#)
- [Verification on page 218](#)

### 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 23 on page 202 shows the topology used in this example.

Figure 23: RSVP-Signaled Point-to-Multipoint LSP



### 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@P2# set protocols ospf traffic-engineering

user@P3# 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 7.7.0/24 {
            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 218](#)
- [Verifying the State of the Point-to-Multipoint LSP on page 219](#)
- [Checking the Forwarding Table on page 219](#)

---

### 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.

```
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 35](#) in the *Junos OS MPLS Applications Configuration Guide*

## Configuring Inter-domain P2MP 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*.

For more information on how to configure link protection, see *Configuring Node Protection or Link Protection for LSPs*.

## 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](#).

---

## 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 [Junos OS Policy Framework Configuration Guide](#).



**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 245](#) and [Configuring BFD for LDP LSPs](#) 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 195](#).

### 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 245](#) and [Configuring BFD for LDP LSPs](#).

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 9

# Miscellaneous MPLS Properties Configuration Guidelines

- [Configuring the Maximum Number of MPLS Labels on page 225](#)
- [Configuring MPLS to Pop the Label on the Ultimate-Hop Router on page 227](#)
- [Configuring Traffic Engineering for LSPs on page 227](#)
- [Enabling Interarea Traffic Engineering on page 230](#)
- [Enabling Inter-AS Traffic Engineering for LSPs on page 231](#)
- [Configuring MPLS to Gather Statistics on page 233](#)
- [Configuring System Log Messages and SNMP Traps for LSPs on page 234](#)
- [Configuring MPLS Firewall Filters and Policers on page 236](#)
- [Configuring MPLS Rewrite Rules on page 243](#)
- [Configuring BFD for MPLS IPv4 LSPs on page 245](#)
- [Pinging LSPs on page 247](#)
- [Tracing MPLS and LSP Packets and Operations on page 249](#)

### 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 5 on page 226 summarizes the label requirements.

**Table 5: 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 31](#)
- Tunneling LDP LSPs in RSVP LSPs Overview
- *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.

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]**



**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 “[Label Description](#)” on page 11 and “[Label Allocation](#)” on page 12.

## 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 230.

The following sections describe how to configure traffic engineering for LSPs:

- [Using LSPs for Both BGP and IGP Traffic Forwarding on page 228](#)
- [Using LSPs for Forwarding in Virtual Private Networks on page 228](#)
- [Using RSVP and LDP Routes for Forwarding but Not Route Selection on page 229](#)
- [Advertising the LSP Metric in Summary LSAs on page 230](#)

## 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-signaled or RSVP-signaled, 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 **bgpp-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-signaled and RSVP-signaled 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 Configuration Guide](#).

---

## 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 196](#).

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 EGBP. Details are provided in the following sections:

- [Inter-AS Traffic Engineering Requirements on page 231](#)
- [Inter-AS Traffic Engineering Limitations on page 232](#)
- [Configuring OSPF Passive TE Mode on page 233](#)

### 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 EGBP is used as the routing protocol between the ASs.
- ASBR information is available inside each AS.
- EGBP 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 233](#).

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 Configuration Guide](#).

## 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 MPLS statistics collection, include the **statistics** statement:

```
statistics {
  auto-bandwidth;
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval seconds;
}
```

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

```

## 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 [Junos OS Network Management Configuration Guide](#).

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 236](#)
- [Examples: Configuring MPLS Firewall Filters on page 237](#)
- [Configuring Policers for LSPs on page 238](#)
- [Example: Configuring an LSP Policer on page 239](#)
- [Configuring Automatic Policers on page 240](#)
- [Writing Different DSCP and EXP Values in MPLS-Tagged IP Packets on page 243](#)

### 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 [Junos OS Policy Framework Configuration Guide](#). For more information about how to configure interfaces, see the [Junos OS Network Interfaces Configuration Guide](#) and the [Junos OS Services Interfaces Configuration Guide](#).

## 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 [Junos OS Policy Framework Configuration Guide](#).

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).

### 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 39](#).

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:** 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 240](#)
- [Configuring Automatic Policers for DiffServ-Aware Traffic Engineering LSPs on page 241](#)
- [Configuring Automatic Policers for Point-to-Multipoint LSPs on page 242](#)
- [Disabling Automatic Policing on an LSP on page 242](#)
- [Example: Configuring Automatic Policing for an LSP on page 242](#)

### 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 240. 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 200](#)). 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 240](#).

### 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 [Junos OS Class of Service Configuration Guide](#). For instructions on how to configure firewall filters, see the [Junos OS Policy Framework Configuration Guide](#).

## 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 [Junos OS Class of Service Configuration Guide](#).

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 243](#)
- [Rewriting MPLS and IPv4 Packet Headers on page 244](#)

### 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 [Junos OS Class of Service Configuration Guide](#).

## 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 [Junos OS Class of Service Configuration Guide](#).

## 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).

For configuration instructions for LDP-signaled LSPs, see *Configuring BFD for LDP LSPs*. 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. 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 seconds;
}
```

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.
- **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 [Junos OS System Basics and Services Command Reference](#).

## 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.

## Pinging LSPs

The following sections describe how to use the **ping mpls** command to confirm LSP functioning.

- [Pinging MPLS LSPs on page 247](#)
- [Pinging Point-to-Multipoint LSPs on page 248](#)
- [Pinging the Endpoint Address of MPLS LSPs on page 248](#)
- [Pinging CCC LSPs on page 248](#)
- [Pinging Layer 3 VPNs on page 248](#)
- [Support for LSP Ping and Traceroute Commands Based on RFC 4379 on page 248](#)

## 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 [Junos OS Routing Protocols and Policies Command Reference](#).



**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 [Junos OS Routing Protocols and Policies Command Reference](#).

## 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 [Junos OS Routing Protocols and Policies Command Reference](#).

## 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 Configuration Guide](#) and the [Junos OS Routing Protocols and Policies Command Reference](#).

## Support for LSP Ping and Traceroute Commands Based on RFC 4379

The Junos OS partially supports LSP **ping** and **traceroute** commands based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. However, the Junos OS only supports this functionality on LSP transit routers. If a **ping** or **traceroute** command is issued from a router that fully supports RFC 4379, it can propagate correctly on routers running the Junos OS.

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 egress 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 Configuration Guide](#).





## PART 3

# Administration

- [Complete MPLS Applications Configuration Statements on page 253](#)
- [MPLS Standards and Support on page 265](#)
- [MPLS Router Configuration Guidelines Reference on page 269](#)
- [DiffServ-Aware Traffic Engineering Configuration Guidelines Reference on page 271](#)
- [Summary of MPLS Configuration Statements on page 281](#)



# Complete MPLS Applications Configuration Statements

- [\[edit logical-systems\] Hierarchy Level on page 253](#)
- [\[edit protocols connections\] Hierarchy Level on page 254](#)
- [\[edit protocols ldp\] Hierarchy Level on page 254](#)
- [\[edit protocols link-management\] Hierarchy Level on page 256](#)
- [\[edit protocols mpls\] Hierarchy Level on page 257](#)
- [\[edit protocols rsvp\] Hierarchy Level on page 261](#)

## [\[edit logical-systems\] 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 [Junos OS Routing Protocols Configuration Guide](#).



**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 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 ];
    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;
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);
}
}

```

## [\[edit protocols link-management\] 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>;
        }
    }
}
}

```

## [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;
    diffserv-te {
      bandwidth-model {
        extended-mam;
        mam;
        rdm;
      }
      te-class-matrix {
        tnumber {
          priority priority;
          traffic-class ctnumber priority priority;
        }
      }
    }
    explicit-null;
    hop-limit number;
    icmp-tunneling;
    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;
description text;
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;
install {
    destination-prefix/prefix-length <active>;
}
ldp-tunneling;
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;
no-cspf;
no-decrement-ttl;
node-link-protection;
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);
}
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>;
    }
}
log-updown {
    no-trap {
        mpls-lsp-traps;
        rfc3812-traps;
    }
    (syslog | no-syslog);
    trap;
    trap-path-down;
    trap-path-up;
}
oam {
    lsp-ping-interval;
    mpls-tp-mode;
    traceoptions;
}
no-cspf;
no-decrement-ttl;
no-propagate-ttl;
optimize-aggressive;
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;
    }
    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;
}

```

## [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;  
}  
}  
}
```



# MPLS Standards and Support

- [Supported MPLS Standards on page 265](#)
- [Link-Layer Support on page 267](#)

## Supported MPLS Standards

---

The 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 [sic] Label Switching (MPLS) Support of Differentiated Services*  
Only E-LSPs are supported.

- RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol [sic] 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 [sic] Label Switched (MPLS) Data Plane Failures*

The traceroute functionality is supported only on transit routers.

- RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*
- Internet draft draft-ietf-bfd-mpls-02.txt, *BFD for MPLS LSPs*
- Internet draft draft-ietf-mpls-rsvp-te-p2mp-01.txt, *Extensions to RSVP-TE for Point to Multipoint TE LSPs* (expires June 2005)
- Internet draft draft-ietf-mpls-soft-preemption-02.txt, *MPLS Traffic Engineering Soft preemption*

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 3063, *MPLS Loop Prevention Mechanism*
- RFC 3208, *PGM Reliable Transport Protocol Specification*

Only the network element is supported.

- RFC 3469, *Framework for Multi-Protocol [sic] 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*

The 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 which 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
- Supported LDP Standards
- Supported RSVP Standards
- Accessing Standards Documents on the Internet



## Link-Layer Support

---

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 Router Configuration Guidelines Reference

- [Minimum MPLS Configuration on page 269](#)

## 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.

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 Labels” on page 11.](#))



## CHAPTER 13

# DiffServ-Aware Traffic Engineering Configuration Guidelines Reference

- DiffServ-Aware Traffic Engineering Standards on page 271
- DiffServ-Aware Traffic Engineering Terminology on page 271
- Configuring Routers for DiffServ-Aware Traffic Engineering on page 272
- Class Type Bandwidth and the LOM on page 277
- LOM Calculation for the MAM and Extended MAM Bandwidth Models on page 277
- LOM Calculation for the Russian Dolls Bandwidth Model on page 277
- Example: LOM Calculation on page 278

### 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**

<b>CAC</b>	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.
<b>Class type</b>	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**

<b>Differentiated Services</b>	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.
<b>Differentiated Services domain</b>	The routers in a network that have Differentiated Services enabled.
<b>DiffServ-aware traffic engineering</b>	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**

<b>MAM</b>	The maximum allocation bandwidth constraint model divides the available bandwidth between the different classes. Sharing of bandwidth between the class types is not allowed.
<b>Multiclass LSP</b>	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**

<b>RDM</b>	The Russian dolls bandwidth constraint model makes efficient use of bandwidth by allowing the class types to share bandwidth.
------------	---

**T**

<b>Traffic engineering class</b>	A paired class type and priority.
<b>Traffic engineering class map</b>	A map between the class types, priorities, and traffic engineering classes. The traffic engineering class mapping must be consistent across the Differentiated Services domain.

---

## 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 274](#)
- [Configuring Traffic Engineering Classes on page 274](#)
- [Configuring Class of Service for DiffServ-Aware Traffic Engineering on page 276](#)

## 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 6 on page 274](#) 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 6: 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



**Table 6: Default Values for the Traffic Engineering Class Matrix (*continued*)**

Traffic Engineering Class	Class Type	Queue	Priority
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 [Junos OS Class of Service Configuration Guide](#).

## 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<sub>c</sub>**—LOM for class type **c**.
- **BC<sub>c</sub>**—Bandwidth constraint for class type **c**.
- **CT<sub>c</sub>**—Class type **c**.
- **TE-Class(i) <--> ( CT<sub>c</sub> , 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:

$$\begin{aligned} \text{Unreserved TE-Class} (i) = \text{LOM}_c \times \text{MIN} [ \\ & [ \text{BC}_c - \text{SUM} ( \text{Normalized} ( \text{CT}_b, q ) ) ] \text{ for } q \leq p \text{ and } c \leq b \leq 7, \\ & \dots \\ & [ \text{BC}_0 - \text{SUM} ( \text{Normalized} ( \text{CT}_b, q ) ) ] \text{ for } q \leq p \text{ and } 0 \leq b \leq 7, \\ & ] \end{aligned}$$

where:

- **LOMc**—LOM for class type **c**.
- **BCc**—Bandwidth constraint for class type **c**.
- **TE-Class(i) <- ->(CTc , 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:

$$\begin{aligned} \text{ct0} &= 8 \times (100 - 10) = 720 \\ \text{ct1} &= 4 \times \min((100-10), (60-10)) = 200 \\ \text{ct2} &= 2 \times \min((100-10), (60-10), (30-10)) = 40 \\ \text{ct3} &= 1 \times \min((100-10), (60-10), (30-10), (10-10)) = 0 \end{aligned}$$

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.



# Summary of MPLS Configuration Statements

## adaptive

---

<b>Syntax</b>	adaptive;
<b>Hierarchy Level</b>	[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> ( <a href="#">primary</a>   <a href="#">secondary</a> ) <i>path-name</i> ], [edit protocols mpls label-switched-path <i>lsp-name</i> ], [edit protocols mpls label-switched-path <i>lsp-name</i> ( <a href="#">primary</a>   <a href="#">secondary</a> ) <i>path-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Default</b>	The configured object is disabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Adaptive LSPs on page 170</a></li> </ul>

## adjust-interval

---

<b>Syntax</b>	<code>adjust-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the bandwidth reallocation interval.
<b>Options</b>	<b><i>seconds</i></b> —Bandwidth reallocation interval, in seconds. <b>Range:</b> 300 through 4,294,967,295 seconds <b>Default:</b> 86,400 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Automatic Bandwidth Allocation Interval on page 157</a></li></ul>

## adjust-threshold

---

<b>Syntax</b>	<code>adjust-threshold <i>percent</i>;</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify how sensitive the automatic bandwidth adjustment for a label-switched path (LSP) is to changes in bandwidth utilization.
<b>Options</b>	<b><i>percent</i></b> —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.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Automatic Bandwidth Adjustment Threshold on page 158</a></li></ul>



## adjust-threshold-overflow-limit

<b>Syntax</b>	adjust-threshold-overflow-limit <i>number</i> ;
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Specify the number of consecutive bandwidth overflow samples before triggering a bandwidth adjustment.
<b>Options</b>	<i>number</i> —Number of consecutive bandwidth overflow samples. <b>Range:</b> 1 through 65,535 <b>Default:</b> This feature is disabled by default.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Limit on Bandwidth Overflow and Underflow Samples on page 159</a></li> </ul>

## adjust-threshold-underflow-limit

<b>Syntax</b>	adjust-threshold-underflow-limit <i>number</i> ;
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.3.
<b>Description</b>	Specify the number of consecutive bandwidth underflow samples before triggering a bandwidth adjustment.
<b>Options</b>	<i>number</i> —Number of consecutive bandwidth underflow samples. <b>Range:</b> 1 through 65,535 <b>Default:</b> This feature is disabled by default.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Limit on Bandwidth Overflow and Underflow Samples on page 159</a></li> </ul>

## admin-down

---

<b>Syntax</b>	admin-down;
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2.
<b>Description</b>	Set a nonpacket GMPLS LSP to the administrative down state. This statement does not affect control path setup or data forwarding for packet LSPs.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Allowing Non-Packet GMPLS LSPs to Establish Paths Through Routers Running the Junos OS</a></li></ul>

## admin-group (for Interfaces)

---

<b>Syntax</b>	admin-group [ <i>group-names</i> ];
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> ], [edit protocols mpls interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define administrative groups for an interface.
<b>Options</b>	<i>group-names</i> —One or more names of groups defined with the <a href="#">admin-groups</a> statement at the [edit protocols mpls] hierarchy level.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Administrative Groups on page 163</a></li><li>• <a href="#">admin-groups on page 289</a></li></ul>

## admin-group (for LSPs)

<b>Syntax</b>	<pre>admin-group {   exclude [ group-names ];   include-all [ group-names ];   include-any [ group-names ]; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls],  [edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i>],  [edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i>  (primary   secondary) <i>path-name</i>],  [edit protocols mpls],  [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i>],  [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> (primary   secondary) <i>path-name</i>]</p>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define the administrative groups to include or exclude an LSP and a path's primary and secondary paths.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.  routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Administrative Groups on page 163</a></li> </ul>

## admin-group-extended

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1.
<b>Description</b>	<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 <b>admin-groups-extended-range</b> 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>
<b>Options</b>	<p><b>apply-groups</b>—Apply the specified administrative groups for the LSP or for the primary and secondary paths.</p> <p><b>apply-groups-except</b>—Exclude the specified administrative groups from the LSP or from the primary and secondary paths.</p> <p><b>exclude</b>—Define the administrative groups to exclude from an LSP or from the primary and secondary paths.</p> <p><b>include-all</b>—Require the LSP to traverse links that include all of the defined administrative groups.</p> <p><b>include-any</b>—Define the administrative groups to include for an LSP for the primary and secondary paths.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Extended Administrative Groups on page 165</a></li> <li>• <a href="#">Configuring Administrative Groups on page 163</a></li> <li>• <a href="#">admin-groups-extended on page 287</a></li> <li>• <a href="#">admin-groups-extended-range on page 288</a></li> </ul>

## admin-groups-extended

---

<b>Syntax</b>	<code>admin-groups-extended <i>group-name</i> {     group-value <i>group-identifier</i>; }</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1.
<b>Description</b>	Specifies the group name and group identifier for an administrative group. The group identifier must be within the range of values specified by the <b>admin-groups-extended-range</b> 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.
<b>Options</b>	<p><b><i>group-name</i></b>—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><b><i>group-value group-identifier</i></b>—The group identifier must be within the range of configurable values, 32 and 4,294,967,295.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Extended Administrative Groups on page 165</a></li> <li>• <a href="#">Configuring Administrative Groups on page 163</a></li> <li>• <a href="#">admin-group-extended on page 286</a></li> <li>• <a href="#">admin-groups-extended-range on page 288</a></li> </ul>

## admin-groups-extended-range

---

<b>Syntax</b>	<pre>admin-groups-extended-range {     maximum <i>maximum-number</i>;     minimum <i>minimum-number</i>; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-options], [edit routing-options]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1.
<b>Description</b>	<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 IGPs (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>
<b>Options</b>	<p><b>maximum</b> <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><b>minimum</b> <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>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Extended Administrative Groups on page 165</a></li><li>• <a href="#">Configuring Administrative Groups on page 163</a></li><li>• <a href="#">admin-group-extended on page 286</a></li></ul>

---

## admin-groups

---

<b>Syntax</b>	<code>admin-groups {     <i>group-name</i> <i>group-value</i>; }</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure administrative groups to implement link coloring of resource classes.
<b>Options</b>	<p><b><i>group-name</i></b>—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><b><i>group-value</i></b>—Value assigned to the group. The names and their corresponding values must be identical across all routers within a single domain.</p> <p><b>Range:</b> 0 through 31</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Administrative Groups on page 163</a></li><li>• <a href="#">admin-group (for Interfaces) on page 284</a></li></ul>

## admin-group-extended

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1.
<b>Description</b>	<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 <b>admin-groups-extended-range</b> 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>
<b>Options</b>	<p><b>apply-groups</b>—Apply the specified administrative groups for the LSP or for the primary and secondary paths.</p> <p><b>apply-groups-except</b>—Exclude the specified administrative groups from the LSP or from the primary and secondary paths.</p> <p><b>exclude</b>—Define the administrative groups to exclude from an LSP or from the primary and secondary paths.</p> <p><b>include-all</b>—Require the LSP to traverse links that include all of the defined administrative groups.</p> <p><b>include-any</b>—Define the administrative groups to include for an LSP for the primary and secondary paths.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Extended Administrative Groups on page 165</a></li> <li>• <a href="#">Configuring Administrative Groups on page 163</a></li> <li>• <a href="#">admin-groups-extended on page 287</a></li> <li>• <a href="#">admin-groups-extended-range on page 288</a></li> </ul>



## admin-groups-extended

<b>Syntax</b>	<code>admin-groups-extended <i>group-name</i> {     group-value <i>group-identifier</i>; }</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.1.
<b>Description</b>	Specifies the group name and group identifier for an administrative group. The group identifier must be within the range of values specified by the <b>admin-groups-extended-range</b> 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.
<b>Options</b>	<p><b><i>group-name</i></b>—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><b><i>group-value group-identifier</i></b>—The group identifier must be within the range of configurable values, 32 and 4,294,967,295.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Extended Administrative Groups on page 165</a></li> <li>• <a href="#">Configuring Administrative Groups on page 163</a></li> <li>• <a href="#">admin-group-extended on page 286</a></li> <li>• <a href="#">admin-groups-extended-range on page 288</a></li> </ul>

## advertisement-hold-time

---

<b>Syntax</b>	<code>advertisement-hold-time seconds;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Do not advertise when the LSP goes from up to down, for a certain period of time known as the hold time.
<b>Options</b>	<b>seconds</b> —Hold time, in seconds. <b>Range:</b> 0 through 65,535 seconds <b>Default:</b> 5 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Damping Advertisement of LSP State Changes on page 178</a></li></ul>

## allow-fragmentation

---

<b>Syntax</b>	<code>allow-fragmentation;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls path-mtu], [edit protocols mpls path-mtu]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Allow IP packets to be fragmented before they are encapsulated in MPLS.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring MTU Signaling in RSVP</a></li></ul>

## always-mark-connection-protection-tlv

<b>Syntax</b>	always-mark-connection-protection-tlv;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-systems-name</i> protocols mpls interface <i>interface-name</i> ], [edit protocols mpls interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.2.
<b>Description</b>	<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 <b>switch-away-lsps</b> statement.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Switching LSPs Away from a Network Node</li> </ul>

## associate-backup-pe-groups

<b>Syntax</b>	associate-backup-pe-groups;
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.0.
<b>Description</b>	<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>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Enabling Point-to-Point LSPs to Monitor Egress PE Routers on page 223</a></li> </ul>

## associate-lsp

---

<b>Syntax</b>	<code>associate-lsp <i>lsp-name</i> {     from <i>from-ip-address</i>; }</code>
<b>Hierarchy Level</b>	<code>[edit protocols mpls label-switched-path <i>lsp-name</i> oam]</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	Configure associated bidirectional label-switched paths (LSPs) on the two ends of an LSP for sending and receiving GAL and G-Ach OAM messages.
<b>Options</b>	<b>from <i>from-ip-address</i></b> —(Optional) Source address for the associated LSP configuration. If omitted, this is derived from the <b>to</b> address of the ingress LSP configuration.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the MPLS Transport Profile for OAM on page 123</a></li></ul>

## auto-bandwidth

---

<b>Syntax</b>	<code>auto-bandwidth {     adjust-interval <i>seconds</i>;     adjust-threshold <i>percent</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>;     monitor-bandwidth; }</code>
<b>Hierarchy Level</b>	<code>[edit protocols mpls label-switched-path <i>lsp-name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Allow an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Automatic Bandwidth Allocation for LSPs on page 155</a></li><li>• <code>request mpls lsp adjust-autobandwidth</code></li></ul>

## auto-policing

---

<b>Syntax</b>	<pre> auto-policing {   class all (drop   loss-priority-high   loss-priority-low);   class <i>ctnumber</i> (drop   loss-priority-high   loss-priority-low); } </pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Enable the automatic policing of all the MPLS LSPs on the router or logical system.
<b>Options</b>	<p><b>class all</b>—Apply the same policer action to all the class types (<b>ct0</b>, <b>ct1</b>, <b>ct2</b>, and <b>ct3</b>).</p> <p><b>class <i>ctnumber</i></b>—Specific class type (<b>ct0</b>, <b>ct1</b>, <b>ct2</b>, or <b>ct3</b>) to which to apply a policer action.</p> <p>Policer actions—You can specify the following policer actions:</p> <p><b>Default:</b> no action</p> <ul style="list-style-type: none"> <li>• <b>drop</b>—Drop all packets.</li> <li>• <b>loss-priority-high</b>—Set the packet loss priority (PLP) to high.</li> <li>• <b>loss-priority-low</b>—Set the PLP to low.</li> </ul>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">policing on page 344</a></li> <li>• <a href="#">Configuring Automatic Policers on page 240v</a></li> </ul>

## backup-pe-group

---

Syntax	<pre>backup-pe-group <i>pe-group-name</i> {     backups [ <i>addresses</i> ];     local-address <i>address</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 9.0.
Description	Configure a backup provider edge (PE) group for ingress PE router redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.
Options	<p><b>backups <i>addresses</i></b>—Specify the address of backup PE routers for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.</p> <p><b>local-address <i>address</i></b>—Specify the address of the local PE router for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.</p> <p><b><i>pe-group-name</i></b>—Specify the name for the group of PE routers that provide ingress PE router redundancy for point-to-multipoint LSPs.</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"><li>• <a href="#">Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs on page 223</a></li></ul>

## bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)

<b>Syntax</b>	<pre> bandwidth <i>bps</i> {     ct0 <i>bps</i>;     ct1 <i>bps</i>;     ct2 <i>bps</i>;     ct3 <i>bps</i>; } </pre>
<b>Hierarchy Level</b>	<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> <b>fast-reroute</b>],  [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> <b>fast-reroute</b>],  [edit protocols mpls label-switched-path <i>lsp-name</i> (primary   secondary) <i>path-name</i>]</p>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<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 <b>ctnumber bandwidth</b> statements to specify the bandwidth to be allocated for each class type.</p>
<b>Options</b>	<p><b>bps</b>—Bandwidth, in bits per second. You can specify this as an integer value. You can also use the abbreviations <b>k</b> (for a thousand), <b>m</b> (for a million), or <b>g</b> (for a billion).</p> <p><b>Range:</b> Any positive integer</p> <p><b>Default:</b> 0 (no bandwidth is reserved)</p> <p><b>ctnumber bps</b>—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 <b>k</b> (for a thousand), <b>m</b> (for a million), or <b>g</b> (for a billion [also called a thousand million]).</p> <p><b>Range:</b> Any positive integer</p> <p><b>Default:</b> 0 (no bandwidth is reserved)</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Fast Reroute on page 144</a></li> <li>• <a href="#">Configuring the Bandwidth Value for LSPs on page 176</a></li> <li>• <a href="#">Configuring Traffic-Engineered LSPs on page 182</a></li> </ul>

- [Configuring Class-Type Bandwidth Constraints for Multiclass LSPs on page 185](#)

## bandwidth (Static LSP)

---

<b>Syntax</b>	<code>bandwidth <i>bps</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> bypass],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</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>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	When configuring a static LSP, specify the traffic rate associated with the LSP.
<b>Options</b>	<b><i>bps</i></b> —Bandwidth, in bits per second. You can specify this as an integer value. You can also use the abbreviations <b>k</b> (for a thousand), <b>m</b> (for a million), or <b>g</b> (for a billion). <b>Range:</b> Any positive integer <b>Default:</b> 0 (no bandwidth is reserved)
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	• <a href="#">Configuring Static LSPs on page 189</a>



## bandwidth-model

---

<b>Syntax</b>	<pre>bandwidth-model {     extended-mam;     mam;     rdm; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls diffserv-te], [edit protocols mpls diffserv-te]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the bandwidth model for differentiated services. Note that you cannot configure both bandwidth models at the same time.
<b>Options</b>	<p><b>extended-mam</b>—The extended maximum allocation model (MAM) is a bandwidth model based on MAM.</p> <p><b>mam</b>—The MAM is defined in RFC 4125, <i>Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering</i>.</p> <p><b>rdm</b>—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>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Bandwidth Model on page 274</a></li> </ul>

## bandwidth-percent

---

<b>Syntax</b>	<code>bandwidth-percent <i>percentage</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> <a href="#">fast-reroute</a> ], [edit protocols mpls label-switched-path <i>lsp-name</i> <a href="#">fast-reroute</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<b><i>percentage</i></b> —The percentage of the protected path's bandwidth that is reserved for the detour path.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Fast Reroute on page 144</a></li><li>• <a href="#">Configuring Fast Reroute for Traffic-Engineered LSPs on page 183</a></li><li>• <a href="#">Configuring Fast Reroute for Multiclass LSPs on page 186</a></li></ul>

## bfd-liveness-detection

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit protocols mpls label-switched-path <i>lsp-name</i> oam], [edit protocols mpls oam]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. <b>failure-action</b> option added in Junos OS Release 8.5.
<b>Description</b>	Enable Bidirectional Forwarding Detection (BFD) for all of the MPLS LSPs or for just a specific LSP.
<b>Options</b>	<p><b>minimum-interval</b>—Minimum transmit and receive interval.  <b>Range:</b> 50 through 255,000 milliseconds  <b>Default:</b> 50</p> <p><b>minimum-receive-interval</b>—Minimum receive interval.  <b>Range:</b> 50 through 255,000 milliseconds  <b>Default:</b> 50</p> <p><b>minimum-transmit-interval</b>—Minimum transmit interval.  <b>Range:</b> 50 through 255,000 milliseconds  <b>Default:</b> 50</p> <p><b>multiplier</b>—Detection time multiplier.  <b>Range:</b> 1 through 255  <b>Default:</b> 3</p> <p>The <b>failure-action</b> statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for MPLS IPv4 LSPs on page 245</a></li> </ul>

## class-of-service

---

<b>Syntax</b>	<code>class-of-service cos-value;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols mpls],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path</code> <code>  <i>lsp-name</i> ingress],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i></code> <code>  (<a href="#">primary</a>   <a href="#">secondary</a>) <i>path-name</i>],</code> <code>[edit protocols mpls],</code> <code>[edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i>],</code> <code>[edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> (<a href="#">primary</a>   <a href="#">secondary</a>) <i>path-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path</code> <code>  <i>lsp-name</i> ingress]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<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>
<b>Options</b>	<p><b>cos-value</b>—CoS value. A higher value typically corresponds to a higher level of service.</p> <p><b>Range:</b> 0 through 7</p> <p><b>Default:</b> 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>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Class of Service for MPLS LSPs on page 167</a></li><li>• <a href="#">Configuring the Ingress Router for Static LSPs on page 189</a></li><li>• <a href="#">Configuring the Intermediate (Transit) and Egress Routers for Static LSPs on page 192</a></li></ul>

## description

---

<b>Syntax</b>	<code>description text;</code>
<b>Hierarchy Level</b>	<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 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 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> <p>[edit protocols mpls static-label-switched-path <i>lsp-name</i> transit <i>incoming-label</i>]</p>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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 <b>show mpls lsp detail</b> command and has no effect on the operation of the LSP.
<b>Options</b>	<b>text</b> —Provide a textual description of the LSP. The description text can be no more than 80 characters in length.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Text Description for LSPs on page 144</a></li> </ul>

## diffserv-te

---

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify properties for differentiated services in traffic engineering.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Routers for DiffServ-Aware Traffic Engineering on page 272</a></li></ul>

## disable

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">interface interface-name</a> ], [edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path lsp-name</a> ], [edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path lsp-name</a> auto-bandwidth], [edit protocols mpls], [edit protocols mpls <a href="#">interface interface-name</a> ], [edit protocols mpls <a href="#">label-switched-path lsp-name</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Disable the functionality of the configured object.
<b>Default</b>	The configured object is enabled (operational) unless explicitly disabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Minimum MPLS Configuration on page 269</a></li> </ul>

## encoding-type

<b>Syntax</b>	encoding-type (ethernet   packet   pdh   sonet-sdh);
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the encoding type of payload carried by the LSP. It can be any of the following: <ul style="list-style-type: none"> <li>• <b>ethernet</b>—Ethernet</li> <li>• <b>packet</b>—Packet</li> <li>• <b>pdh</b>—Plesiochronous digital hierarchy (PDH)</li> <li>• <b>sonet-sdh</b>—SONET/SDH</li> </ul>
<b>Default</b>	packet
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Encoding Type</a></li> </ul>

## exclude (for Administrative Groups)

---

<b>Syntax</b>	<code>exclude [ <i>group-names</i> ];</code>
<b>Hierarchy Level</b>	<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) path-name 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) path-name admin-group]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define the administrative groups to exclude for an LSP or for a path's primary and secondary paths.
<b>Options</b>	<i>group-names</i> —Names of one or more groups defined with the <a href="#">admin-groups</a> statement.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Administrative Groups on page 163</a></li></ul>

## exclude (for Fast Reroute)

---

<b>Syntax</b>	<code>(exclude [ <i>group-names</i> ]   no-exclude);</code>
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Control exclusion of administrative groups: <ul style="list-style-type: none"><li>• <b>exclude</b>—Define the administrative groups to exclude for fast reroute.</li><li>• <b>no-exclude</b>—Disable administrative group exclusion.</li></ul>
<b>Options</b>	<i>group-names</i> —Names of one or more groups defined with the <a href="#">admin-groups</a> statement.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Fast Reroute on page 144</a></li><li>• <a href="#">admin-groups on page 289</a></li></ul>



## exclude-srlg

<b>Syntax</b>	exclude-srlg;
<b>Hierarchy Level</b>	[edit protocols mpls], [edit logical-systems logical-system-name protocols mpls], [edit protocols mpls label-switched-path <i>path-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>path-name</i> ], [edit 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], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>destination</i> ], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>destination</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	<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>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Excluding SRLG Links Completely for the Secondary LSP on page 78</a></li> </ul>

## expand-loose-hop

---

<b>Syntax</b>	expand-loose-hop;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. Point-to-multipoint LSP support introduced in Junos OS Release 11.2.
<b>Description</b>	<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>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Enabling Interarea Traffic Engineering on page 230</a></li><li>• <a href="#">Configuring Inter-domain P2MP LSPs on page 219</a></li></ul>

## explicit-null

---

<b>Syntax</b>	explicit-null;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Advertise label 0 to the egress router of an LSP.
<b>Default</b>	If you do not include the <b>explicit-null</b> statement in the MPLS configuration, label 3 (implicit null) is advertised.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring RSVP to Pop the Label on the Ultimate-Hop Router</a></li></ul>

## failure-action

---

<b>Syntax</b>	<pre>failure-action {     make-before-break teardown-timeout <i>seconds</i>;     teardown; }</pre>
<b>Hierarchy Level</b>	<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>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4.
<b>Description</b>	<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>
<b>Options</b>	<p><b>make-before-break</b>—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><b>teardown</b>—When a BFD session fails for an RSVP LSP, the associated LSP path is taken down and resigaled immediately.</p> <p><b>teardown-timeout <i>seconds</i></b>—When you configure the <b>make-before-break</b> option, you can specify a time in seconds for the <b>teardown-timeout</b> option. At the end of the time specified, the associated RSVP LSP is automatically torn down and resigaled.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a Failure Action for the BFD Session on an RSVP LSP on page 246</a></li> </ul>

## fast-reroute

---

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	The statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Fast Reroute on page 144</a></li></ul>

## fate-sharing

<b>Syntax</b>	<pre>fate-sharing {   group <i>group-name</i> {     cost <i>value</i>;     from <i>address</i> &lt;to <i>address</i>&gt;;   } }</pre>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.3 for the QFX Series.
<b>Description</b>	Specify groups of objects that share characteristics resulting in backup paths to be used if primary paths fail. All objects are treated as /32 host addresses. You specify one or more objects within a group. The objects can be LAN interfaces, router IDs, or point-to-point links. The sequence is insignificant.
<b>Options</b>	<p><b>cost <i>value</i></b>—Cost assigned to the group.  <b>Range:</b> 1 through 65,535  <b>Default:</b> 1</p> <p><b>from <i>address</i></b>—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 <b>from</b> addresses in the group.</p> <p><b>group <i>group-name</i></b>—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><b>to <i>address</i></b>—(Optional) Address of egress router. For point-to-point link objects, you must specify both a <b>from</b> and a <b>to</b> address.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Alternate Backup Paths Using Fate Sharing on page 51</a></li> </ul>

## from

---

<b>Syntax</b>	<code>from address;</code>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>Specify the source address to use for the LSP.</p> <p>The address you specify does not affect the outgoing interface used by the LSP.</p>
<b>Default</b>	If you do not include this statement, the software automatically selects the loopback interface as the address.
<b>Options</b>	<i>address</i> —IP address.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Ingress Router Address for LSPs on page 139</a></li></ul>

## gpipd

<b>Syntax</b>	<code>gpipd (ethernet   hdlc   ipv4   pos-scrambling-crc-16   pos-no-scrambling-crc-16   pos-scrambling-crc-32   pos-no-scrambling-crc-32   ppp);</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>pos-scrambling-crc-16</b> , <b>pos-no-scrambling-crc-16</b> , <b>pos-scrambling-crc-32</b> , and <b>pos-no-scrambling-crc-32</b> options added in Junos OS Release 8.0.
<b>Description</b>	Specify the type of payload carried by the LSP. It can be any of the following: <ul style="list-style-type: none"> <li>• <b>ethernet</b>—Ethernet (GPID value: 33)</li> <li>• <b>hdlc</b>—High-level Data Link Control (HDLC) (GPID value: 44)</li> <li>• <b>ipv4</b>—IP version 4 (GPID value: 0x0800)</li> <li>• <b>pos-no-scrambling-crc-16</b>—for interoperability with other vendors' equipment (GPID value: 29)</li> <li>• <b>pos-no-scrambling-crc-32</b>—for interoperability with other vendors' equipment (GPID value: 30)</li> <li>• <b>pos-scrambling-crc-16</b>—for interoperability with other vendors' equipment (GPID value: 31)</li> <li>• <b>pos-scrambling-crc-32</b>—for interoperability with other vendors' equipment (GPID value: 32)</li> <li>• <b>ppp</b>—Point-to-Point Protocol (PPP) (GPID value: 50)</li> </ul>
<b>Default</b>	<code>ipv4</code>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• Configuring MPLS LSPs for GMPLS</li> </ul>

## hop-limit

---

<b>Syntax</b>	<code>hop-limit <i>number</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols mpls],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls <i>label-switched-path lsp-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i></code> <code>  <i>fast-reroute</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i></code> <code>  (<i>primary</i>   <i>secondary</i>) <i>path-name</i>],</code> <code>[edit protocols mpls],</code> <code>[edit protocols mpls <i>label-switched-path lsp-name</i>],</code> <code>[edit protocols mpls label-switched-path <i>lsp-name</i> <i>fast-reroute</i>],</code> <code>[edit protocols mpls label-switched-path <i>lsp-name</i> (<i>primary</i>   <i>secondary</i>) <i>path-name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>For an LSP, specify the maximum number of routers that the LSP can traverse, including the ingress and egress routers.</p> <p>For fast reroute, how many more routers a detour is allowed to traverse compared with the LSP itself. 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.</p>
<b>Options</b>	<p><i>number</i>—Maximum number of hops.</p> <p><b>Range:</b> 2 through 255 (for an LSP); 0 through 255 (for fast reroute)</p> <p><b>Default:</b> 255 (for an LSP); 6 (for fast reroute)</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Fast Reroute on page 144</a></li><li>• <a href="#">Limiting the Number of Hops in LSPs on page 176</a></li></ul>



## icmp-tunneling

<b>Syntax</b>	icmp-tunneling;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Enable ICMP message tunneling for MPLS LSPs. This feature helps you to trace the route path and debug LSPs.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring ICMP Message Tunneling for MPLS on page 69</a></li> </ul>

## include-all (for Administrative Groups)

<b>Syntax</b>	include-all [ <i>group-names</i> ];
<b>Hierarchy Level</b>	[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 <b>label-switched-path</b> <i>lsp-name</i> ( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> admin-group], [edit protocols mpls label-switched-path <i>lsp-name</i> admin-group], [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> admin-group]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Require the LSP to traverse links that include all of the defined administrative groups.
<b>Options</b>	<b><i>group-names</i></b> —One or more names of groups defined with the <b>admin-groups</b> statement.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Administrative Groups on page 163</a></li> <li>• <a href="#">admin-groups on page 289</a></li> </ul>

## include-all (for Fast Reroute)

---

Syntax	(include-all [ <i>group-names</i> ]   no-include-all);
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	Control inclusion of administrative groups: <ul style="list-style-type: none"><li>• <b>include-all</b>—Define the administrative groups that must all be included for fast reroute.</li><li>• <b>no-include-all</b>—Disable administrative group inclusion.</li></ul>
Options	<b>group-names</b> —One or more names of groups defined with the <b>admin-groups</b> 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"><li>• <a href="#">Configuring Fast Reroute on page 144</a></li></ul>

## include-any (for Administrative Groups)

---

Syntax	include-any [ <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 <b>label-switched-path</b> <i>lsp-name</i> ( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> admin-group], [edit protocols mpls label-switched-path <i>lsp-name</i> admin-group], [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> admin-group]
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	<b>group-names</b> —One or more names of groups defined with the <b>admin-groups</b> 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"><li>• <a href="#">Configuring Administrative Groups on page 163</a></li></ul>

## include-any (for Fast Reroute)

---

<b>Syntax</b>	(include-any [ <i>group-names</i> ]   no-include-any);
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Control inclusion of administrative groups: <ul style="list-style-type: none"><li>• <b>include-any</b>—Define the administrative groups to include for fast reroute.</li><li>• <b>no-include-any</b>—Disable administrative group inclusion.</li></ul>
<b>Options</b>	<b><i>group-names</i></b> —One or more names of groups defined with the <b>admin-groups</b> statement.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Fast Reroute on page 144</a></li></ul>

## ingress

---

Syntax	<pre>ingress {   bandwidth <i>bps</i>;   class-of-service <i>cos-value</i>;   description <i>string</i>;   install {     destination-prefix &lt;active&gt;;   }   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	[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.
Description	Configure an ingress LSR for a static 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"><li>• <a href="#">Configuring Static LSPs on page 189</a></li></ul>

## install

---

<b>Syntax</b>	install { <i>destination-prefix</i> <active>; }
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Associate one or more prefixes with an LSP. When the LSP is up, all the prefixes are installed as entries into the <b>inet.3</b> routing table.
<b>Options</b>	<b>active</b> —(Optional) Install the route into the <b>inet.0</b> routing table. This allows you to issue a <b>ping</b> or <b>traceroute</b> command on this address.  <b><i>destination-prefix</i></b> —Address to associate with the LSP.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Adding LSP-Related Routes to the inet.3 Routing Table on page 146</a></li> </ul>

## interface

---

Syntax	interface ( <i>interface-name</i>   all) { disable; admin-group [ <i>group-names</i> ]; srlg <i>srlg-name</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	Enable MPLS on one or more interfaces.
Options	<b>interface-name</b> —Name of the interface on which to configure MPLS. To configure all interfaces, specify <b>all</b> . For details about specifying interfaces, see the <a href="#">Junos OS Network Interfaces Configuration Guide</a> .  <b>srlg srlg-name</b> —Name of the SRLG to associate with an interface.  The remaining 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"><li>• <a href="#">Configuring the Intermediate (Transit) and Egress Routers for Static LSPs on page 192</a></li><li>• <a href="#">Example: Configuring SRLG on page 69</a></li></ul>

## 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.
Description	Allow IPv6 routes to be resolved over an MPLS network by converting all routes stored in the <b>inet.3</b> routing table to IPv4-mapped IPv6 addresses and then copying them into the <b>inet6.3</b> routing table. This routing table can be used to resolve next hops for both <b>inet6</b> and <b>inet6-vpn</b> 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"><li>• <a href="#">Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks on page 60</a></li></ul>

## label-switched-path

```

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;
        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>;
        }
        ldp-tunneling;
        link-protection;
        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);
        }
        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>;
}
}

```

<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure an LSP to use in dynamic MPLS. When configuring an LSP, you must specify the address of the egress router in the <b>to</b> statement. All remaining statements are optional.
<b>Options</b>	<p><b><i>lsp-name</i></b>—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>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Minimum MPLS Configuration on page 269</a></li> <li>• <a href="#">LSP Configuration Overview on page 136</a></li> </ul>

## ldp-tunneling

<b>Syntax</b>	ldp-tunneling;
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Enable the LSP to be used for LDP tunneling.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Miscellaneous LDP Properties</a></li> </ul>

## least-fill

See [random](#)

## link-protection (Dynamic LSPs)

---

<b>Syntax</b>	link-protection;
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<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 <b>link-protection</b> 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>
<b>Default</b>	Link protection is disabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Link Protection for Point-to-Multipoint LSPs on page 221</a></li><li>• Configuring Node Protection or Link Protection for LSPs</li><li>• link-protection (RSVP)</li></ul>

## link-protection (Static LSPs)

<b>Syntax</b>	link-protection bypass-name <i>name</i> ;
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	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.
<b>Default</b>	Link protection is disabled.
<b>Options</b>	bypass-name <i>name</i> —Bypass LSP name.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Static LSPs on page 189</a></li> <li>• Example: Configuring Point-to-Multipoint LSPs with Static Routes</li> </ul>

## log-updown

---

<b>Syntax</b>	<pre>log-updown {   no-trap {     mpls-lsp-traps;     rfc3812-traps;   }   (syslog   no-syslog);   trap;   trap-path-down;   trap-path-up; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. The <b>mpls-lsp-traps</b> and <b>rfc-3812-traps</b> options added in Junos OS Release 9.0.
<b>Description</b>	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.
<b>Default</b>	There is no default behavior for this statement. If you do not specify the options, the configuration cannot be committed.
<b>Options</b>	<b>no-syslog</b> —Do not log a message to the system log file.  <b>no-trap</b> —Do not send an SNMP trap.  <b>syslog</b> —Log a message to the system log file.  <b>trap</b> —Send an SNMP trap.  <b>trap-path-down</b> —Send an SNMP trap when an LSP path goes down.  <b>trap-path-up</b> —Send an SNMP trap when an LSP path comes up.  The <b>no-trap</b> statement is explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring System Log Messages and SNMP Traps for LSPs on page 234</a></li><li>• <a href="#">Junos OS Network Management Configuration Guide</a></li><li>• <a href="#">no-trap on page 337</a></li><li>• <a href="#">traceoptions on page 371</a></li></ul>

## **lsp-attributes**

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>pos-scrambling-crc-16</b> , <b>pos-no-scrambling-crc-16</b> , <b>pos-scrambling-crc-32</b> , and <b>pos-no-scrambling-crc-32</b> options added in Junos OS Release 8.0.
<b>Description</b>	Define the parameters signaled during LSP setup. These usually determine the nature of the resource (label) allocated for the LSP.  The options are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Configuring MPLS LSPs for GMPLS</li> </ul>

## **maximum-bandwidth**

<b>Syntax</b>	maximum-bandwidth <i>bps</i> ;
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the maximum amount of bandwidth in bits per second (bps).
<b>Options</b>	<i>bps</i> —Maximum amount of bandwidth.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 158</a></li> </ul>

## maximum-labels

---

<b>Syntax</b>	<code>maximum-labels <i>maximum-labels</i>;</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	<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"><li>• MX Series 3D Universal Edge Router</li><li>• M120 Multiservice Edge Router</li><li>• M320 Multiservice Edge Router with Enhanced III FPCs</li><li>• M7i Multiservice Edge Router and M10i Multiservice Edge Router with Enhanced Compact Forwarding Engine Board (CFEB-E)</li><li>• T640, T1600, TX Matrix, and TX Matrix Plus routers with Enhanced Scaling FPC1, Enhanced Scaling FP2, Enhanced Scaling FPC3, and Enhanced Scaling FPC4</li></ul>
<b>Options</b>	<p><i>maximum-labels</i>—Maximum number of labels.</p> <p><b>Range:</b> 3 through 5</p> <p><b>Default:</b> 3</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Maximum Number of MPLS Labels on page 225</a></li><li>• <a href="#">Junos OS VPNs Configuration Guide</a></li></ul>

## metric

<b>Syntax</b>	<code>metric <i>metric</i>;</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<i>metric</i> —LSP metric value. <b>Default:</b> No metric assigned (dynamic) <b>Range:</b> 1 through 16,777,215
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Static LSP Metrics on page 148</a></li> </ul>

## minimum-bandwidth

<b>Syntax</b>	<code>minimum-bandwidth <i>bps</i>;</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Set the minimum bandwidth in bps for an LSP with automatic bandwidth allocation enabled.
<b>Options</b>	<i>bps</i> —Minimum bandwidth for the LSP.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 158</a></li> </ul>

## monitor-bandwidth

---

<b>Syntax</b>	monitor-bandwidth;
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Passive Bandwidth Utilization Monitoring on page 161</a></li></ul>

## most-fill

---

See [random](#)

## mpls

---

<b>Syntax</b>	mpls { ... }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols], [edit protocols]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Enable MPLS on the router.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Minimum MPLS Configuration on page 269</a></li></ul>



## mpls-tp-mode

---

<b>Syntax</b>	<code>mpls-tp-mode;</code>
<b>Hierarchy Level</b>	[edit protocols mpls label-switched-path <i>lsp-name</i> oam], [edit protocols mpls oam]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	<p>Enable GAL or G-Ach OAM operation without IP encapsulation on a label-switched path (LSP).</p> <p>Include this statement at the <b>[edit protocols mpls oam]</b> 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 <b>[edit protocols mpls label-switched-path <i>lsp-name</i> oam]</b> hierarchy level to enable GAL and G-Ach OAM operation without IP encapsulation on a specific LSP.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring the MPLS Transport Profile for OAM on page 123</a></li> </ul>

## mtu-signaling

---

<b>Syntax</b>	<code>mtu-signaling;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls path-mtu rsvp], [edit protocols mpls path-mtu rsvp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Enable MTU signaling in RSVP.
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring MTU Signaling in RSVP</a></li> </ul>

## next-hop

---

<b>Syntax</b>	<code>next-hop (address   interface-name   address/interface-name);</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name bypass],</code> <code>[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name ingress],</code> <code>[edit logical-systems logical-system-name protocols mpls static-label-switched-path lsp-name transit incoming-label],</code> <code>[edit protocols mpls static-label-switched-path lsp-name bypass],</code> <code>[edit protocols mpls static-label-switched-path lsp-name ingress],</code> <code>[edit protocols mpls static-label-switched-path lsp-name transit incoming-label]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	IP address of the next hop to the destination, specified as the IP address of the next hop, the interface name (for point-to-point interfaces only), or the <b>address/interface-name</b> to specify an IP address on an operational interface.
<b>Options</b>	<b>address</b> —IP address of the next-hop router.  <b>interface-name</b> —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.
<b>Required Privilege Level</b>	<b>routing</b> —To view this statement in the configuration. <b>routing-control</b> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Ingress Router for Static LSPs on page 189</a></li><li>• <a href="#">Configuring the Intermediate (Transit) and Egress Routers for Static LSPs on page 192</a></li></ul>

## no-cspf

<b>Syntax</b>	no-cspf;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> ], [edit protocols mpls], [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ], [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<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"> <li>• Interior gateway protocol (IGP) topology database</li> <li>• Link utilization information from extensions in the IGP link-state database</li> <li>• Administrative group information from extensions in the IGP link-state database</li> <li>• LSP requirements, including bandwidth, hop count, and administrative group</li> </ul> <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>
<b>Default</b>	Constrained-path LSP computation enabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Disabling Constrained-Path LSP Computation on page 162</a></li> <li>• <a href="#">Configuring Explicit-Path LSPs on page 196</a></li> </ul>

## no-decrement-ttl

---

<b>Syntax</b>	no-decrement-ttl;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ( <a href="#">primary</a>   <a href="#">secondary</a> ) <i>path-name</i> ], [edit protocols mpls], [edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ], [edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ( <a href="#">primary</a>   <a href="#">secondary</a> ) <i>path-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Default</b>	Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Disabling Normal TTL Decrementing on page 153</a></li><li>• <a href="#">no-propagate-ttl on page 336</a></li></ul>

## no-install-to-address

---

<b>Syntax</b>	no-install-to-address;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <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 <a href="#">label-switched-path</a> <i>lsp-name</i> ], [edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.4.
<b>Description</b>	Prevent the egress router address configured using the <b>to</b> statement from being installed into the <b>inet.3</b> and <b>inet.0</b> routing tables.
<b>Default</b>	The egress router address for an LSP is installed into the <b>inet.3</b> and <b>inet.0</b> routing tables.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Preventing the Addition of Egress Router Addresses to Routing Tables on page 140</a></li> <li>• <a href="#">to on page 370</a></li> </ul>

## no-propagate-ttl

---

<b>Syntax</b>	no-propagate-ttl;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Default</b>	Normal TTL decrementing enabled; the TTL field value is decremented by 1 as the packet passes through each label-switched router in the LSP.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Disabling Normal TTL Decrementing on page 153</a></li><li>• Example: Disabling Normal TTL Decrementing in a VRF Routing Instance (on Layer 3 VPNs Configuration Guide or in the <i>Junos VPNs Configuration Guide</i>)</li><li>• <a href="#">no-decrement-ttl on page 334</a></li></ul>

## no-trap

---

<b>Syntax</b>	no-trap { mpls-lsp-traps; rfc-3812-traps; }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls log-updown], [edit protocols mpls log-updown]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. The <b>mpls-lsp-traps</b> and <b>rfc-3812-traps</b> options added in Junos OS Release 9.0.
<b>Description</b>	Prevent the transmission of SNMP traps.
<b>Options</b>	<p><b>mpls-lsp-traps</b>—Block the MPLS LSP traps defined in the <b>jnx-mpls.mib</b>, but allows the <b>rfc3812.mib</b> traps.</p> <p><b>rfc-3812-traps</b>—Block the traps defined in the <b>rfc3812.mib</b>, but allows the MPLS LSP traps defined in the <b>jnx-mpls.mib</b>.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring System Log Messages and SNMP Traps for LSPs on page 234</a></li> <li>• <a href="#">Junos OS Network Management Configuration Guide</a></li> <li>• <a href="#">traceoptions on page 371</a></li> </ul>

## node-protection (Static LSP)

---

<b>Syntax</b>	<code>node-protection bypass-name <i>name</i> next-next-label <i>label</i>;</code>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced in JUNOS Release 10.1.
<b>Description</b>	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.
<b>Default</b>	Node protection is disabled.
<b>Options</b>	<code>bypass-name <i>name</i></code> —Bypass LSP name. <code>next-next-label <i>label</i></code> —Bypass LSP name.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Static LSPs on page 189</a></li></ul>



## oam

<b>Syntax</b>	<pre>oam {   bfd-liveness-detection{     failure-action teardown;     minimum-interval <i>milliseconds</i>;     minimum-receive-interval <i>milliseconds</i>;     minimum-transmit-interval <i>milliseconds</i>;     multiplier <i>detection-time-multiplier</i>;   }   lsp-ping-interval <i>seconds</i>; }</pre>
<b>Hierarchy Level</b>	[edit protocols mpls], [edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ] [edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> primary <i>path-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.6. <b>lsp-ping-interval</b> option introduced in Junos OS Release 9.4.
<b>Description</b>	Enable Operation, Administration, and Maintenance (OAM) for RSVP-signaled LSPs.
<b>Options</b>	<p><b>lsp-ping-interval <i>seconds</i></b>—Specify the duration of the LSP ping interval in seconds. To issue a ping on an RSVP-signaled LSP, use the <b>ping mpls rsvp</b> command.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring BFD for MPLS IPv4 LSPs on page 245</a></li> </ul>

## optimize-aggressive

---

<b>Syntax</b>	optimize-aggressive;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Default</b>	Aggressive optimization is disabled.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Optimizing Signaled LSPs on page 172</a></li></ul>

## optimize-timer

<b>Syntax</b>	<code>optimize-timer <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls],          [edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i>],          [edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i>            (<b>primary</b>   <b>secondary</b>) <i>path-name</i>],          [edit protocols mpls],          [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i>],          [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> (<b>primary</b>   <b>secondary</b>) <i>path-name</i>]</p>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<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 option is useful only on LSPs for which constrained-path computation is enabled; that is, for which the <b>no-cspf</b> statement is not configured.</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>
<b>Default</b>	The optimize timer is disabled.
<b>Options</b>	<p><b><i>seconds</i></b>—Length of the optimize timer, in seconds.</p> <p><b>Range:</b> 0 through 65,535 seconds</p> <p><b>Default:</b> 0 seconds (the optimize timer is disabled)</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Optimizing Signaled LSPs on page 172</a></li> </ul>

## p2mp

---

<b>Syntax</b>	<code>p2mp p2mp-lsp-name;</code>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<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.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Primary Point-to-Multipoint LSP on page 200</a></li></ul>

## p2mp-lsp-next-hop

---

<b>Syntax</b>	<code>p2mp-lsp-next-hop point-to-multipoint-lsp-next-hop;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-options static route <i>route-name</i> ], [edit routing-options static route <i>route-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the name of the point-to-multipoint LSP to be used as a next hop for the static route.
<b>Options</b>	<i>point-to-multipoint-lsp-next-hop</i> —Name of the point-to-multipoint LSP.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Static Unicast Routes for Point-to-Multipoint LSPs on page 195</a></li></ul>

## path

<b>Syntax</b>	<code>path <i>path-name</i> {     (<i>address</i>   <i>hostname</i>) &lt;strict   loose&gt;; }</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Create a named path and optionally specify the sequence of explicit routers that form the path.  You must include this statement when configuring explicit LSPs.
<b>Options</b>	<p><b><i>address</i></b>—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 <b>loose</b>. 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><b>Default:</b> If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.</p> <p><b><i>hostname</i></b>—See <b><i>address</i></b>.</p> <p><b>Default:</b> If you do not specify any routers explicitly, no routing limitations are imposed on the LSP.</p> <p><b><i>loose</i></b>—(Optional) Indicate that the next address in the <b>path</b> statement is a loose link. This means that the LSP can traverse through other routers before reaching this router.</p> <p><b>Default:</b> <b>strict</b></p> <p><b><i>path-name</i></b>—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><b><i>strict</i></b>—(Optional) Indicate that the LSP must go to the next address specified in the <b>path</b> statement without traversing other nodes. This is the default.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Creating Named Paths on page 50</a></li> </ul>

## path-mtu

---

<b>Syntax</b>	<pre>path-mtu {     allow-fragmentation;     rsvp {         mtu-signaling;     } }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure MTU options for MPLS paths, including packet fragmentation and MTU signaling.  The remaining statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>Configuring MTU Signaling in RSVP</li></ul>

## policing

---

<b>Syntax</b>	<pre>policing {     filter <i>filter-name</i>;     no-auto-policing; }</pre>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the policing filter for the LSP.
<b>Options</b>	<b>filter <i>filter-name</i></b> —Specify the name of the policing filter.  <b>no-auto-policing</b> —Disable automatic policing on this LSP.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li><a href="#">Configuring Policers for LSPs on page 238</a></li><li><a href="#">auto-policing on page 295</a></li></ul>

---

## pop

---

<b>Syntax</b>	pop;
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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).
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Intermediate (Transit) and Egress Routers for Static LSPs on page 192</a></li><li>• <a href="#">swap on page 366</a></li></ul>

## preference

---

<b>Syntax</b>	<code>preference <i>preference</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols mpls],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls <i>label-switched-path lsp-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols mpls <i>label-switched-path lsp-name</i></code> <code>  (<i>primary</i>   <i>secondary</i>) <i>path-name</i>],</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 protocols mpls],</code> <code>[edit protocols mpls <i>label-switched-path lsp-name</i>],</code> <code>[edit protocols mpls <i>label-switched-path lsp-name</i> (<i>primary</i>   <i>secondary</i>) <i>path-name</i>],</code> <code>[edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<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>
<b>Options</b>	<p><b><i>preference</i></b>—Preference to assign to the route. A route with a lower preference value is preferred.</p> <p><b>Range:</b> 1 through 255</p> <p><b>Default:</b> 5 for static MPLS LSPs, 7 for RSVP MPLS LSPs, 9 for LDP MPLS LSPs</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Preference Values for LSPs on page 167</a></li><li>• <a href="#">Configuring the Ingress Router for Static LSPs on page 189</a></li><li>• <a href="#">Configuring the Intermediate (Transit) and Egress Routers for Static LSPs on page 192</a></li></ul>



## primary

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<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 <b>[edit mpls label-switched-path <i>lsp-name</i>]</b> hierarchy level).</p>
<b>Options</b>	<p><b><i>path-name</i></b>—Name of a path that you created with the <b>path</b> statement.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">Configuring Primary and Secondary LSPs on page 141</a></li> </ul>

## priority

---

<b>Syntax</b>	<code>priority setup-priority reservation-priority;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> ], [edit protocols mpls], [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ], [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<p><b>reservation-priority</b>—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><b>Range:</b> 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p><b>Default:</b> 0 (Once the session is set up, no other session can preempt it.)</p> <p><b>setup-priority</b>—Setup priority.</p> <p><b>Range:</b> 0 through 7, where 0 is the highest and 7 is the lowest priority.</p> <p><b>Default:</b> 7 (The session cannot preempt any existing sessions.)</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Priority and Preemption for LSPs on page 171</a></li></ul>

## protection-revert-time

---

<b>Syntax</b>	<code>protection-revert-time <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.1.
<b>Description</b>	<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 <b>protection-revert-time</b> 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>
<b>Options</b>	<p><b><i>seconds</i></b>—Time in seconds.</p> <p><b>Range:</b> 0 through 65,535 seconds</p> <p><b>Default:</b> 5 seconds</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Static LSPs on page 189</a></li> </ul>

## push

---

<b>Syntax</b>	<code>push out-label;</code>
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<b>out-label</b> —Manually assigned outgoing label value. <b>Range:</b> 0 through 1,048,575.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">pop on page 345</a></li><li>• <a href="#">swap on page 366</a></li><li>• <a href="#">Configuring the Ingress Router for Static LSPs on page 189</a></li></ul>

## random

---

<b>Syntax</b>	(random   least-fill   most-fill);
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<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"> <li>• <b>least-fill</b>—Prefer the path with the most available bandwidth (with the largest minimum available bandwidth ratio).</li> <li>• <b>most-fill</b>—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.</li> <li>• <b>random</b>—Choose the path at random.</li> </ul>
<b>Default</b>	random
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring CSPF Tie Breaking on page 149</a></li> </ul>

## record

<b>Syntax</b>	(record   no-record);
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> ( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> ], [edit protocols mpls], [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ], [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> ( <b>primary</b>   <b>secondary</b> ) <i>path-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Default</b>	Record routes.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Disabling Path Route Recording on page 167</a></li> </ul>

## retry-limit

<b>Syntax</b>	retry-limit <i>number</i> ;
<b>Hierarchy Level</b>	[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> ],
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Options</b>	<i>number</i> —Maximum number of tries to establish the primary path. <b>Range:</b> 0 through 10,000 <b>Default:</b> 0 (The ingress node never stops trying to establish the primary path.)
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Connection Between Ingress and Egress Routers on page 147</a></li> </ul>

## retry-timer

---

<b>Syntax</b>	<code>retry-timer <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Amount of time the ingress router waits between attempts to establish the primary path.
<b>Options</b>	<b><i>seconds</i></b> —Amount of time between attempts to connect to the primary path. <b>Range:</b> 1 through 600 seconds <b>Default:</b> 30 seconds
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Connection Between Ingress and Egress Routers on page 147</a></li></ul>

## revert-timer

---

<b>Syntax</b>	<code>revert-timer <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. BFD behavior modified in Junos OS Release 9.0.
<b>Description</b>	<p>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.</p> <p>If you have configured BFD on the LSP, the Junos OS waits until the BFD session is restored before starting the revert timer counter.</p> <p>If you have configured a value of 0 seconds for the <b>revert-timer</b> 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.</p>
<b>Options</b>	<p><b><i>seconds</i></b>—Time in seconds.</p> <p><b>Range:</b> 0 through 65,535 seconds</p> <p><b>Default:</b> 60 seconds</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Revert Timer for LSPs on page 142</a></li></ul>



## rpf-check-policy

---

<b>Syntax</b>	<code>rpf-check-policy <i>policy</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-options multicast], [edit routing-options multicast]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.0.
<b>Description</b>	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.
<b>Options</b>	<i>policy</i> —Name of the RPF check routing policy.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring a Multicast RPF Check Policy for Point-to-Multipoint LSPs on page 222</a></li></ul>

## rsvp-error-hold-time

---

<b>Syntax</b>	<code>rsvp-error-hold-time seconds;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<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>
<b>Options</b>	<p><b>seconds</b>—Amount of time MPLS retains RSVP PathErr messages and considers them for CSPF computations.</p> <p><b>Range:</b> 0 through 240 seconds</p> <p><b>Default:</b> 25 seconds</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages on page 57</a></li></ul>

## secondary

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<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>
<b>Options</b>	<p><b><i>path-name</i></b>—Name of a path that you created with the <b>path</b> statement.</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Primary and Secondary LSPs on page 141</a></li> </ul>

## select

<b>Syntax</b>	<code>select (manual   unconditional);</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i> (primary   secondary) <i>path-name</i> ], [edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i> (primary   secondary) <i>path-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the conditions under which the path is selected to carry traffic. The <b>manual</b> and <b>unconditional</b> options are mutually exclusive.
<b>Options</b>	<p><b>manual</b>—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).</p> <p><b>unconditional</b>—The path is always selected for carrying traffic, even if it is currently down or degraded (receiving errors).</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Specifying the Conditions for Path Selection on page 143</a></li> </ul>

## signal-bandwidth

<b>Syntax</b>	<code>signal-bandwidth <i>type</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> <i>lsp-attributes</i> ], [edit protocols mpls label-switched-path <i>lsp-name</i> <i>lsp-attributes</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the bandwidth encoding of the signal used for path computation and admission control.
<b>Options</b>	<b>type</b> —Configure the type of bandwidth encoding used on the LSP. It can be any of the following values: <b>10gigether</b> , <b>ds1</b> , <b>ds3</b> , <b>e1</b> , <b>e3</b> , <b>ethernet</b> , <b>fastether</b> , <b>gigether</b> , <b>stm-1</b> , <b>stm-4</b> , <b>stm-16</b> , <b>stm-64</b> , <b>stm-256</b> , <b>sts-1</b> , <b>vt1-5</b> , or <b>vt2</b> .
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring MPLS LSPs for GMPLS</a></li> </ul>

## smart-optimize-timer

<b>Syntax</b>	<code>smart-optimize-timer <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	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.
<b>Default</b>	The smart optimization timer is disabled by default.
<b>Options</b>	<p><b><i>seconds</i></b>—(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><b>Range:</b> 0 through 65,535 seconds</p> <p><b>Default:</b> 180 seconds</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring the Smart Optimize Timer on page 175</a></li> <li>• <a href="#">Optimizing Signaled LSPs on page 172</a></li> <li>• <a href="#">optimize-aggressive on page 340</a></li> <li>• <a href="#">optimize-timer on page 341</a></li> </ul>

## soft-preemption

---

<b>Syntax</b>	soft-preemption;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ], [edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Attempt to establish a new path for a preempted LSP before tearing it down.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring MPLS Soft Preemption on page 154</a></li></ul>

## srlg

---

<b>Syntax</b>	srlg { <i>srlg-name</i> { srlg-cost <i>srlg-cost</i> ; srlg-value <i>srlg-value</i> ; } }
<b>Hierarchy Level</b>	[edit routing-options], [edit logical-systems <i>logical-system-name</i> routing-options] [edit protocols mpls interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	Configure Shared Risk Link Group (SRLG) parameters.
<b>Options</b>	<b>srlg-cost <i>srlg-cost</i></b> —Specify a cost for the SRLG ranging from 1 through 65535. <b>srlg-value <i>srlg-value</i></b> —Specify a Group ID for the SRLG ranging from 1 through 4294967295.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring SRLG on page 69</a></li></ul>

## srlg-cost

---

<b>Syntax</b>	<code>srlg-cost <i>srlg-cost</i>;</code>
<b>Hierarchy Level</b>	[edit routing-options srlg], [edit logical-systems <i>logical-system-name</i> routing-options srlg]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	Specify a cost for the Shared Risk Link Group (SRLG) ranging from 1 through <b>65535</b> .
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring SRLG on page 69</a></li> </ul>

## srlg-value

---

<b>Syntax</b>	<code>srlg-value <i>srlg-value</i>;</code>
<b>Hierarchy Level</b>	[edit routing-options srlg], [edit logical-systems <i>logical-system-name</i> routing-options srlg]
<b>Release Information</b>	Statement introduced in Junos OS Release 11.4.
<b>Description</b>	Specify a Group ID for the Shared Risk Link Group (SRLG) ranging from 1 through <b>4294967295</b> .
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring SRLG on page 69</a></li> </ul>

## standby

---

<b>Syntax</b>	standby;
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Have the path remain up at all times to provide instant switchover if connectivity problems occur.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Hot Standby of Secondary Paths on page 177</a></li></ul>



## 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.

**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 189](#)

## statistics

<b>Syntax</b>	<pre>statistics {   auto-bandwidth;   file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;   interval <i>seconds</i>; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Enable MPLS statistics collection and reporting.
<b>Options</b>	<p><b>auto-bandwidth</b>—Collect statistics related to automatic bandwidth.</p> <p><b>file <i>filename</i></b>—(Optional) Name of the file to receive the output. We recommend that you place MPLS tracing output in the file <b>mpls-stat</b> in the <b>/var/log</b> directory.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b><i>file</i></b> reaches its maximum size, it is renamed <b><i>file.0</i></b>, then <b><i>file.1</i></b>, and so on, until the maximum number of files is reached. Then, the oldest file is overwritten.</p> <p><b>Range:</b> 2 or more</p> <p><b>Default:</b> 2 files</p> <p>If you specify a maximum number of files, you also must specify a maximum file size with the <b>size</b> option.</p> <p><b>interval <i>seconds</i></b>—Interval at which to periodically collect statistics.</p> <p><b>Range:</b> 1 through 65,535</p> <p><b>Default:</b> 300 seconds</p> <p><b>no-world-readable</b>—(Optional) Prevent users from reading the log file.</p> <p><b>size <i>size</i></b>—(Optional) Maximum size of each file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a file named <b><i>file</i></b> reaches this size, it is renamed <b><i>file.0</i></b>. When the <b><i>file</i></b> again reaches its maximum size, <b><i>file.0</i></b> is renamed <b><i>file.1</i></b> and <b><i>file</i></b> is renamed <b><i>file.0</i></b>. This renaming scheme continues until the maximum number of files is reached. Then the oldest trace file is overwritten.</p> <p><b>Syntax:</b> Syntax: <b><i>xk</i></b> to specify KB, <b><i>xm</i></b> to specify MB, or <b><i>xg</i></b> to specify GB</p> <p><b>Range:</b> 10 KB through the maximum file size supported on your system</p> <p><b>Default:</b> 1 MB</p> <p>If you specify a maximum file size, you also must specify a maximum number of files with the <b>files</b> option.</p> <p><b>world-readable</b>—(Optional) Enable users to read the log file.</p>

**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 233](#)
- [Configuring MPLS Statistics for Automatic Bandwidth Allocation on page 156](#)

---

## swap

---

**Syntax** `swap out-label;`

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols mpls static-label-switched-path *lsp-name* transit *incoming-label*],  
[edit protocols mpls static-label-switched-path *lsp-name* transit *incoming-label*]

**Release Information** Statement introduced before Junos OS Release 7.4.

**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**

- [pop on page 345](#)
- [push on page 350](#)
- [Configuring the Intermediate \(Transit\) and Egress Routers for Static LSPs on page 192](#)

## switch-away-lsps

---

<b>Syntax</b>	switch-away-lsps;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-systems-name</i> protocols mpls interface <i>interface-name</i> ], [edit protocols mpls interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.2.
<b>Description</b>	<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 <b>always-mark-connection-protection-tlv</b> statement.</p> <p>The <b>always-mark-connection-protection-tlv</b> 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 <b>switch-away-lsps</b> statement, traffic is switched to the bypass LSP.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>Switching LSPs Away from a Network Node</li> </ul>

## switching-type

---

<b>Syntax</b>	switching-type (fiber   lambda   psc-1   tdm);
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>Specify the switching method for the LSP. The switching method can be one of the following values:</p> <ul style="list-style-type: none"><li>• <b>fiber</b>—Fiber switching</li><li>• <b>lambda</b>—Lambda switching</li><li>• <b>psc-1</b>—Packet switching</li><li>• <b>tdm</b>—Time-division multiplexing (TDM) switching</li></ul>
<b>Default</b>	psc-1
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• Configuring MPLS LSPs for GMPLS</li></ul>

## te-class-matrix

<b>Syntax</b>	<pre>te-class-matrix {   tenumber {     priority <i>priority</i>;     traffic-class {       ctnumber priority <i>priority</i>;     }   } }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls diffserv-te], [edit protocols mpls diffserv-te]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the traffic engineering class matrix for a multiclass LSP or a DiffServ-aware traffic engineering LSP.
<b>Default</b>	<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>
<b>Options</b>	<p><b>ctnumber</b>—Specify the number of the class type. It can be one of four values: <b>ct0</b>, <b>ct1</b>, <b>ct2</b>, or <b>ct3</b>.</p> <p><b>priority <i>priority</i></b>—Specify the priority of the class type. It can be one of eight values from 0 through 7.</p> <p><b>tenumber</b>—Specify the number of the traffic engineering class. It can be one of eight values: <b>te0</b>, <b>te1</b>, <b>te2</b>, <b>te3</b>, <b>te4</b>, <b>te5</b>, <b>te6</b>, or <b>te7</b>. You must configure the traffic engineering classes in order, starting with <b>te0</b>.</p> <p><b>traffic-class</b>—Specify the traffic class for the traffic engineering class.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Traffic Engineering Classes on page 274</a></li> </ul>

to

---

<b>Syntax</b>	<code>to address;</code>
<b>Hierarchy Level</b>	[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> bypass], [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> bypass], [edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the egress router of a dynamic LSP.
<b>Options</b>	<i>address</i> —Address of the egress router.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring the Egress Router Address for LSPs on page 139</a></li></ul>



## traceoptions

<b>Syntax</b>	<pre>traceoptions {     file <i>filename</i> &lt;files <i>number</i>&gt; &lt;size <i>size</i>&gt; &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i>; }</pre>
<b>Hierarchy Level</b>	[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> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>Configure MPLS tracing options at the protocol level or for a label-switched path.</p> <p>To specify more than one tracing operation, include multiple <b>flag</b> statements.</p>
<b>Default</b>	The default MPLS protocol-level tracing options are inherited from the routing protocols <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.
<b>Options</b>	<p><b><i>filename</i></b>—Name of the file to receive the output of the tracing operation. All files are placed in the directory <b>/var/log</b>. We recommend that you place MPLS tracing output in the file <b>mpls-log</b>.</p> <p><b><i>files number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b><i>trace-file</i></b> reaches its maximum size, it is renamed <b><i>trace-file.0</i></b>, then <b><i>trace-file.1</i></b>, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.</p> <p><b>Range:</b> 2 through 1000</p> <p><b>Default:</b> 2 files</p> <p>If you specify a maximum number of files, you must also include the <b>size</b> statement to specify the maximum file size.</p> <p><b><i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <p><b>MPLS Tracing Flags</b></p> <ul style="list-style-type: none"> <li>• <b>all</b>—Trace all operations</li> <li>• <b>connection</b>—All circuit cross-connect (CCC) activity</li> <li>• <b>connection-detail</b>—Detailed CCC activity</li> <li>• <b>cspf</b>—CSPF computations</li> <li>• <b>cspf-link</b>—Links visited during CSPF computations</li> <li>• <b>cspf-node</b>—Nodes visited during CSPF computations</li> <li>• <b>error</b>—MPLS error packets</li> </ul>

- **graceful-restart**—Trace MPLS graceful restart 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
- **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.

<b>Required Privilege Level</b>	routing and trace—To view this statement in the configuration.
	routing-control and trace-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tracing MPLS and LSP Packets and Operations on page 249</a></li></ul>

## traffic-engineering

<b>Syntax</b>	traffic-engineering (bgp   bgp-igp   bgp-igp-both-ribs   mpls-forwarding);
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	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.
<b>Default</b>	bgp
<b>Options</b>	<p><b>bgp</b>—On BGP destinations only. Ingress routes are installed in the <b>inet.3</b> routing table.</p> <p><b>bgp-igp</b>—On both BGP and IGP destinations. Ingress routes are installed in the <b>inet.0</b> routing table. If IGP shortcuts are enabled, the shortcut routes are automatically installed in the <b>inet.0</b> routing table.</p> <p><b>bgp-igp-both-ribs</b>—On both BGP and IGP destinations. Ingress routes are installed in the <b>inet.0</b> and <b>inet.3</b> routing tables. This option is used to support VPNs.</p> <p><b>mpls-forwarding</b>—On both BGP and IGP destinations. Use ingress routes for forwarding only, not for routing.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Traffic Engineering for LSPs on page 227</a></li> <li>• Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)</li> </ul>

## transit-lsp-association

---

<b>Syntax</b>	<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>
<b>Hierarchy Level</b>	[edit protocols mpls]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.1.
<b>Description</b>	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.
<b>Options</b>	<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>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring the MPLS Transport Profile for OAM on page 123</a></li></ul>

## PART 4

# Index

- [Index on page 377](#)



# Index

## Symbols

#, comments in configuration statements.....	xx
( ), in syntax descriptions.....	xx
< >, in syntax descriptions.....	xx
[ ], in configuration statements.....	xx
{ }, in configuration statements.....	xx
(pipe), in syntax descriptions.....	xx

## A

adaptive rerouting.....	170, 281
adaptive statement.....	281
usage guidelines.....	170
addresses	
associating with LSPs.....	146, 319
egress router address.....	139, 370
ingress router.....	139, 312
adjust-interval statement.....	282
usage guidelines.....	157
adjust-threshold statement.....	282
usage guidelines.....	158
adjust-threshold-overflow-limit statement.....	283
usage guidelines.....	159
adjust-threshold-underflow-limit statement.....	283
usage guidelines.....	159
admin-down statement.....	284
admin-group statement	
LSPs.....	285
MPLS interfaces.....	284
admin-group-extended statement.....	286, 290
usage guidelines.....	165
admin-groups statement.....	289
usage guidelines.....	163
admin-groups-extended statement.....	287, 291
usage guidelines.....	165
admin-groups-extended-range statement.....	288
usage guidelines.....	165
administrative groups	
admin-groups statement.....	289
configuration.....	163
exclude statement.....	306
extended.....	165

fast reroute.....	144
include-all statement.....	315
include-any statement.....	316
advertisement-hold-time statement.....	292
usage guidelines.....	178
aggregated interfaces.....	25
all (tracing flag).....	371
allocation of labels.....	12
allow-fragmentation statement.....	292
always-mark-connection-protection-tlv	
statement.....	293
associate-backup-pe-groups statement.....	293
usage guidelines.....	223
associate-lsp	
usage guidelines.....	123
associate-lsp statement	
MPLS-TP.....	294
auto-bandwidth statement.....	294
usage guidelines.....	157
auto-policing statement.....	295
usage guidelines.....	240
automatic bandwidth allocation.....	35, 294
bandwidth monitoring.....	161
LSPs.....	155
manually trigger.....	161
threshold.....	158
automatic policers	
LSP bandwidth, changing.....	241
LSPs.....	240
overview.....	240
point-to-multipoint LSPs.....	242

## B

backup paths.....	19
backup-pe-group statement.....	296
bandwidth	
automatic allocation, LSPs.....	155
LSP paths.....	176
bandwidth model.....	274
bandwidth oversubscription	
overview.....	43
bandwidth statement	
fast reroute.....	297
usage guidelines.....	144
LSPs	
usage guidelines.....	182
MPLS	
usage guidelines.....	192

multiclass LSPs.....	297	constrained-path LSPs	
usage guidelines.....	185	computation	
signaled LSPs.....	297	CSPF algorithm.....	15
usage guidelines.....	176	disabling.....	162, 333
static LSPs.....	298	overview.....	15
usage guidelines (ingress router).....	189	example configuration.....	54, 55
bandwidth-model statement.....	299	overview.....	15
usage guidelines.....	274	scope.....	15
bandwidth-percent statement.....	300	conventions	
usage guidelines.....	183, 186	text and syntax.....	xix
BFD		CoS.....	274
fast reroute.....	245	Differentiated Services.....	180
revert timer.....	354	CoS values.....	168
RSVP LSPs.....	245, 246	cross-connect, circuit See CCC	
bfd-liveness-detection statement		cspf (tracing flag).....	371
RSVP LSPs.....	301	CSPF algorithm	
usage guidelines.....	245	fate-sharing.....	53
BGP		offline path computation.....	6, 18
destinations.....	26	online path computation.....	15
braces, in configuration statements.....	xx	disabling.....	162, 333
brackets		overview.....	6
angle, in syntax descriptions.....	xx	cspf-link (tracing flag).....	371
square, in configuration statements.....	xx	cspf-node (tracing flag).....	371
branch LSPs.....	200	curly braces, in configuration statements.....	xx
<b>C</b>		customer support.....	xxi
CCC		contacting JTAC.....	xxi
ping CCC LSPs.....	248	<b>D</b>	
class types		damping	
bandwidth subscription.....	179	LSP transitions.....	178
class-of-service statement		description statement	
ingress routers.....	302	MPLS.....	303
usage guidelines.....	189	usage guidelines.....	144, 192
signaled LSPs.....	302	static LSPs	
usage guidelines.....	168	usage guidelines (ingress router).....	189
static LSPs.....	302	detours See fast reroute	
colored links.....	144, 163, 289	Differentiated Services	
comments, in configuration statements.....	xx	bandwidth model.....	274
connection (tracing flag).....	371	extended MAM.....	274
connection-detail (tracing flag).....	371	interface bandwidth constraints.....	180
connections statement		LSPs.....	40
complete hierarchy under.....	254	MAM.....	274
Constrained Shortest Path First algorithm See CSPF		RDM.....	274
algorithm		diffserv-te statement.....	304
		usage guidelines.....	272
		disable statement	
		MPLS.....	305
		usage guidelines.....	162



documentation  
     comments on.....xx

DSCP  
     MPLS-tagged packets.....243

dynamic LSP metric.....148

**E**

egress routers  
     example configuration.....194  
     overview.....14  
     signaled LSPs.....139  
     static LSPs.....192, 320

empty paths.....343

encoding-type statement.....305

error (tracing flag)  
     MPLS.....371

exclude statement  
     administrative groups  
         usage guidelines.....163  
     fast reroute  
         usage guidelines.....144

exclude-srlg  
     usage guidelines.....78, 83, 103

exclude-srlg statement.....307

EXP and IP precedence bits.....244

EXP bits.....12, 168, 169, 236  
     DSCP values.....243  
     rewrite.....243

EXP rewrite rule.....169

expand-loose-hop statement.....308  
     usage guidelines.....230

experimental bits *See* EXP bits

Explicit Null label.....11

Explicit Route object.....7

explicit routes.....6

explicit-null statement  
     MPLS.....308  
         usage guidelines.....227  
     RSVP.....308

explicit-path LSPs  
     computation, disabling.....162, 333  
     configuring.....196  
     example configuration.....54  
     overview.....15  
     scope.....15

extended administrative groups.....165

extended MAM.....274, 299

**F**

failed LSPs  
     fast reroute.....30, 144, 297, 310  
     standby secondary paths.....30

failure-action statement  
     RSVP LSPs.....309  
         usage guidelines.....246

family mpls statement  
     usage guidelines.....150

fast reroute  
     BFD.....245  
     configuring.....310  
     detours.....31  
     multiclass LSPs.....186  
     overview.....30, 144  
     path optimization.....146  
     path optimization overview.....34  
     PFE fast reroute.....144  
     soft preemption.....154  
     traffic-engineered LSPs.....183

fast-reroute statement.....310  
     RSVP  
         usage guidelines.....146

fate-sharing  
     CSPF algorithm.....53  
     example configuration.....53  
     overview.....19  
     signaled LSPs.....51, 311

fate-sharing statement.....311  
     usage guidelines.....51

font conventions.....xix

forwarding *See* MPLS

forwarding next hop.....28

from statement  
     MPLS.....312  
         usage guidelines.....139

**G**

gpip statement.....313

graceful restart  
     point-to-multipoint LSPs.....221

graceful-restart (tracing flag).....372

GRE tunnels.....59

groups  
     administrative.....144, 163, 289

**H**

headers, MPLS and IPv4.....244

hold priority.....171

hold time	
signaled LSPs.....	178, 292
hop-limit statement.....	314
usage guidelines.....	176
host routes.....	26, 146
hot-standby state.....	177
<b>I</b>	
ICMP	
message tunneling, MPLS.....	69
icmp-tunneling statement.....	315
usage guidelines.....	69
IEEE 802.p rewrite rule.....	169
IGP	
advertising LSPs.....	24
destinations.....	28
shortcuts	
enabling.....	22
LSP metrics.....	148
overview.....	21
qualified LSPs.....	22
routing tables.....	23
uses of.....	23
Implicit Null label.....	11
include statement	
fast reroute	
usage guidelines.....	144
include-all statement	
administrative groups	
usage guidelines.....	163
include-any statement	
administrative groups	
usage guidelines.....	163
inet.0 routing table	
IGP shortcuts.....	23
MPLS.....	29
inet.3 routing table	
IGP shortcuts.....	23
MPLS.....	29
routes, installing.....	146
information distribution, traffic engineering.....	5
ingress routers	
address configuration.....	139, 312
configuring for static LSPs.....	189
example configurations.....	191
overview.....	14
path connection retry	
information.....	147, 352, 353
ingress statement	
static LSP.....	318
ingress static LSPs.....	189
install statement	
MPLS.....	319
usage guidelines.....	146
static LSPs	
usage guidelines (ingress router).....	189
inter-area traffic engineering.....	230
Inter-domain point-to-multipoint LSPs.....	219
interface statement	
static LSPs.....	320
interfaces	
aggregated.....	25
interior gateway protocol See IGP	
intermediate routers	
configuring for static LSPs.....	192, 320
example configurations.....	193
intraregion LSPs.....	23
IP packets over aggregated interfaces.....	25
IPv4 Explicit Null label.....	11
IPv6	
Implicit Null label.....	11
tunneling over MPLS.....	60
ipv6-tunneling statement.....	320
<b>L</b>	
Label Distribution Protocol See LDP	
Label object.....	7
Label Request object.....	7
label-switched paths See LSPs	
label-switched-path statement	
MPLS.....	321
usage guidelines.....	136
MPLS with RSVP.....	321
labels	
allocation.....	12
numerical ranges.....	11
operations.....	13
overview.....	10
properties.....	189
reserved labels.....	11
stacks.....	12
swapping.....	4
values.....	11
LDP	
tunneling through RSVP LSPs.....	323
ldp statement	
complete hierarchy under.....	254

- 
- ldp-tunneling statement.....323
  - least-fill statement.....351
    - usage guidelines.....149
  - least-fill tie-breaking rule.....18, 149, 351
  - link attributes considered by CSPF algorithm.....15
  - link coloring.....144, 163, 289
  - link protection
    - soft preemption.....154
  - link-layer protocols.....267
  - link-management statement
    - complete hierarchy under.....256
  - link-protection statement
    - MPLS
      - usage guidelines.....192, 221
    - signaled LSPs.....324
    - static LSPs.....325
    - usage guidelines (ingress router).....189
  - load balancing
    - MPLS LSPs.....150
  - log-updown statement
    - MPLS.....326
    - usage guidelines.....234
  - logical-router See logical-system
  - logical-routers See logical-systems
  - loose explicit routes.....6, 196
  - LSP metric.....230
  - lsp-attributes statement.....327
  - lsp-ping-interval statement
    - RSVP LSPs.....339
  - lsping (tracing flag).....372
  - LSPs
    - adaptive rerouting.....170, 281
    - administrative groups
      - admin-groups statement.....289
      - configuring.....163
      - fast reroute.....144
    - advertising in IGP.....24
    - associating addresses.....146, 319
    - attributes considered by CSPF algorithm.....15
    - automatic bandwidth allocation.....294
    - automatic policers.....240
    - bandwidth
      - maximum bounds.....158
      - minimum bounds.....158
    - BFD configuration.....245
    - configuration statements.....136, 321
    - configuring.....182
    - constrained-path See constrained-path LSPs
    - CoS values.....168
    - creating.....136
    - damping LSP transitions.....178
    - description, textual.....144
    - differentiated service aware.....40
    - egress routers.....139, 192, 194, 320
    - explicit-path See explicit-path LSPs
    - failure of.....30
    - fast reroute.....30, 144, 297, 310
    - fate-sharing.....19, 51, 311
    - forwarding next hops
      - selecting.....28
    - hold time.....178, 292
    - hop limit.....176
    - host routes.....26
    - IGP shortcuts.....21
    - ingress routers.....139, 312
    - intermediate routers.....192, 320
    - intraregion LSPs.....23
    - load balancing.....150
    - metrics.....148, 329
    - MPLS routers, configuring.....57
    - named paths.....50, 343
    - OAM configuration.....245
    - overview.....4, 10
    - packet traversal.....5, 14
    - path
      - bandwidth.....176
      - calculation.....3
      - connection retry
        - information.....147, 352, 353
      - length.....176, 314
      - smart optimize timer.....175
    - pings.....247
      - ping interval, RSVP.....246
    - policing.....238
    - preemption.....171, 348
    - preference levels.....167, 346
    - primary.....141, 347
    - priorities.....171, 348
    - recording routes.....167
    - reoptimization.....172, 340, 341
    - router functions.....14
    - routing options.....7
    - RSVP-signaled.....15
    - scope of.....15
    - secondary.....141, 357
    - signaled See signaled LSPs
    - soft preemption.....154
    - standby secondary paths.....30

standby state.....	177, 362	IGP and BGP destinations.....	28
static See static LSPs		Implicit Null label.....	227
text description.....	303	IPv4 packet headers.....	244
tie-breaking rules.....	18, 149, 351	IPv6.....	60
traffic engineering, configuring.....	227	label range.....	11
TTL decrementing, disabling.....	153, 334, 336	link-layer protocols supported.....	267
tunneling through RSVP LSPs.....	323	load balancing.....	150
		LSPs See LSPs	
<b>M</b>		OAM.....	245
MAM.....	274, 299	overview.....	10
manuals		packets over aggregated interfaces.....	25
comments on.....	xx	ping	
maximum-bandwidth statement.....	327	Layer 3 VPNs.....	248
usage guidelines.....	158	LSP end points.....	248
maximum-labels statement.....	328	LSPs.....	247
messages		routing tables.....	29
MPLS system log.....	234, 326	RSVP See RSVP	
RSVP PathErr.....	57	signaled LSPs See signaled LSPs	
metric statement		smart optimize timer.....	175
MPLS.....	329	SNMP traps.....	234, 326
usage guidelines.....	148	soft preemption.....	154
static LSPs		standby secondary paths.....	30
usage guidelines (ingress router).....	189	static.....	189, 320
metrics		static LSPs See static LSPs	
dynamic LSP metric.....	148	statistics.....	156
static LSP metric.....	148, 329	statistics output.....	233
minimum-bandwidth statement.....	329	supported software standards.....	265
usage guidelines.....	158	system log messages.....	234, 326
monitor-bandwidth statement.....	330	tracing protocol operations.....	249, 371
usage guidelines.....	161	traffic engineering.....	230
most-fill statement.....	351	overview.....	10
usage guidelines.....	149	traffic protection.....	30
most-fill tie-breaking rule.....	18, 149, 351	traffic statistics.....	233, 365
MPLS.....	3	tunneling	
aggregated interfaces.....	25	IPv6.....	60
automatic bandwidth allocation.....	294	ultimate-hop popping.....	227, 308
backbones, packet traversal.....	5, 14	See also LDP, LSPs, RSVP, traffic engineering	
BFD.....	245, 246	database	
BGP destinations.....	26	mpls statement	
configuring.....	269	Layer 2 switching cross-connect.....	330
CoS values.....	168	MPLS.....	330
DSCP and EXP values.....	243	complete hierarchy under.....	257
EXP bits.....	12, 168, 169, 236	usage guidelines.....	269
Explicit Null label.....	227	mpls transport profile oam	
extended administrative groups.....	165	overview.....	122
fast reroute.....	30, 144, 297, 310	mpls-tp-mode	
firewall filter.....	236	usage guidelines.....	123
GRE tunnels.....	59	mpls-tp-mode statement	
ICMP, tunneling.....	69	MPLS-TP.....	331

mpls.0 routing table.....	29
mtu-signaling statement.....	331
multicast	
RPF check policy.....	222
multiclass LSPs	
bandwidth subscription.....	179
configuring.....	185
fast reroute.....	186
multiple push (label operation).....	13

## N

named paths	
empty paths.....	343
example configuration.....	51
overview.....	50
next hops	
selecting.....	28
next-hop statement.....	332
MPLS	
usage guidelines.....	192
static LSPs	
usage guidelines (ingress router).....	189
no-cspf statement.....	333
usage guidelines.....	162
no-decrement-ttl statement.....	334
no-install-to-address statement.....	335
static LSPs	
usage guidelines (ingress router).....	189
usage guidelines.....	140
no-p2mp-sublsp statement	
usage guidelines.....	224
no-propagate-ttl statement.....	336
no-record statement.....	352
usage guidelines.....	167
no-trap statement.....	337
usage guidelines.....	234
no-world-readable option to traceoptions	
statement	
MPLS.....	372
node protection	
soft preemption.....	154
node-protection statement	
MPLS	
usage guidelines.....	192
static LSPs.....	338
usage guidelines (ingress router).....	189
nsr-synchronization (tracing flag).....	372
nsr-synchronization-detail (tracing flag).....	372

## O

oam statement	
RSVP LSPs.....	339
usage guidelines.....	245
OAM, MPLS.....	245
offline path calculation.....	6, 18
operations on labels.....	13
optimize-aggressive statement.....	340
usage guidelines.....	172
optimize-timer statement	
MPLS.....	341
usage guidelines.....	172
optimizing LSPs.....	172, 340, 341
OSPF	
inter-area traffic engineering.....	230
LSP metric advertisement.....	230

## P

p2mp statement.....	342
usage guidelines.....	200
p2mp-lsp-next-hop statement.....	342
RSVP	
usage guidelines.....	201
usage guidelines.....	195
packet forwarding component	
traffic engineering.....	4
packet headers, MPLS and IPv4.....	244
packet traversal on LSPs.....	5, 14
parentheses, in syntax descriptions.....	xx
path	
bandwidth, LSP.....	176
calculation	
constrained-path computation.....	162, 333
CSPF algorithm.....	6, 15
offline path computation.....	6, 18
routing options.....	7
tie-breaking rules.....	18, 149, 351
connection retry information.....	147, 352, 353
length, LSP.....	176, 314
selection component, traffic engineering.....	6
path optimization	
fast reroute.....	146
path selection.....	143
path statement	
MPLS.....	343
path-mtu statement.....	344
PathErr messages.....	57
PFE fast reroute.....	144

ping	
Layer 3 VPNs.....	248
LSP end point.....	248
LSPs.....	247
point-to-multipoint LSP.....	248
PLP bit.....	168
point-to-multipoint LSPs	
automatic policers.....	242
branch LSPs.....	200
dynamic.....	201
static.....	201
configuration.....	199
graceful restart.....	37, 221
GRES.....	37
inter-domain.....	219
link protection.....	221
overview.....	35
RPF check policy.....	222
static routes.....	195
with RSVP signaling.....	201
policing.....	238, 240
policing filter statement	
usage guidelines.....	238
policing statement.....	344
static LSPs	
usage guidelines (ingress router).....	189
usage guidelines.....	240
pop (label operation).....	13
pop statement	
MPLS.....	345
usage guidelines.....	192
preemption	
LSPs.....	154
signaled LSPs.....	171, 348
preference levels	
signaled LSPs.....	167, 346
static LSPs.....	189
preference statement	
signaled LSPs.....	346
usage guidelines.....	167
static LSPs.....	346
usage guidelines (ingress router).....	189
primary LSPs.....	141, 347
primary paths	
revert timer.....	142
revert timer, BFD.....	354
selection.....	143
primary statement	
MPLS.....	347
usage guidelines.....	141
priorities	
signaled LSPs.....	171, 348
priority statement	
MPLS.....	348
usage guidelines.....	171
protection-revert-time statement.....	349
push (label operation).....	13
push statement	
MPLS.....	350
static LSPs	
usage guidelines (ingress router).....	189
<b>R</b>	
random statement.....	351
usage guidelines.....	149
random tie-breaking rule.....	18, 149, 351
RDM.....	274, 299
Record Route object.....	167
record statement.....	352
usage guidelines.....	167
recording routes.....	167
reoptimizing LSPs.....	172, 340, 341
rerouting LSPs	
adaptive rerouting.....	170, 281
fast reroute.....	30, 144, 297, 310
reserved labels.....	11
reserving network resources See RSVP	
resource classes.....	144, 163
Resource Reservation Protocol See RSVP	
retry information.....	147, 352, 353
retry-limit statement.....	352
usage guidelines.....	147
retry-timer statement.....	353
usage guidelines.....	147
revert-timer statement.....	354
usage guidelines.....	142
rewrite rules.....	169
IEEE 802.p and MPLS CoS.....	169
MPLS and VPNs.....	244
route preferences	
signaled LSPs.....	167, 346
Router Alert label.....	11
routers	
egress See egress routers	
ingress See ingress routers	
label operations.....	13

LSP functions.....	14
transit.....	14
routes	
recording.....	167
route preferences.....	167, 346
routing options, traffic engineering.....	7
routing tables	
host routes, installing.....	146, 319
IGP shortcuts.....	23
inet.0.....	23, 29
inet.3.....	23, 29, 146
MPLS.....	29
mpls.0.....	29
rpf-check-policy statement.....	355
configuration guidelines.....	222
RSVP.....	3
BFD.....	245, 246
Differentiated Services.....	180
for point-to-multipoint LSPs.....	201
PathErr messages.....	57
signaled LSPs.....	15
signaling extensions.....	7
tunneling LDP LSPs through RSVP LSPs.....	323
<i>See also</i> LDP	
rsvp statement	
complete hierarchy under.....	261
rsvp-error-hold-time statement.....	356
usage guidelines.....	57
Russian dolls bandwidth model.....	277

## S

scope of LSPs.....	15
secondary	
LSPs.....	141, 177, 357
paths.....	30
revert timer.....	142
selection.....	143
secondary statement.....	357
usage guidelines.....	141
select statement.....	358
usage guidelines.....	143
setup priority, signaled LSPs.....	171
shared risk link group	
overview.....	20
signal-bandwidth statement.....	358
signaled LSPs.....	343
adaptive rerouting.....	170, 281
administrative groups	
admin-groups statement.....	289
configuring.....	163
fast reroute.....	144
associating addresses.....	146, 319
configuration statements.....	136, 321
constrained-path computation	
disabling.....	162, 333
CoS values.....	168
creating.....	136
damping LSP transitions.....	178
egress router address.....	139, 370
fast reroute.....	144, 297, 310
fate-sharing.....	51, 311
hold time.....	178, 292
ingress router address.....	139, 312
metrics.....	148, 329
MPLS routers, configuring.....	57
named paths.....	50
path	
bandwidth.....	176
connection retry information.....	147, 352
length.....	176, 314
preemption.....	171, 348
preference levels.....	167, 346
primary.....	141, 347
priorities.....	171, 348
recording routes.....	167
reoptimization.....	172, 340, 341
RSVP <i>See</i> RSVP	
secondary.....	141, 357
standby state.....	177, 362
tie-breaking rules.....	18, 149, 351
TTL decrementing.....	153, 334, 336
signaling component, traffic engineering.....	7
signaling extensions, RSVP.....	7
smart-optimize-timer statement.....	359
usage guidelines.....	175
SNMP traps	
MPLS.....	234, 326
soft-preemption statement	
MPLS.....	360
usage guidelines.....	154
RSVP	
usage guidelines.....	154
special labels.....	11

srlg.....	20	swap statement	
excluding, common links, secondary		MPLS.....	366
path.....	78, 83, 103	usage guidelines.....	192
overview.....	20	switch-away-lsps statement.....	367
usage guidelines.....	69	switching-type statement.....	368
srlg statement.....	360	syntax conventions.....	xix
srlg-cost		system log messages	
usage guidelines.....	69	MPLS.....	234, 326
srlg-cost statement.....	361		
srlg-value		<b>T</b>	
usage guidelines.....	69	te-class-matrix statement.....	369
srlg-value statement.....	361	usage guidelines.....	274
stacked labels.....	12	technical support	
standby secondary paths.....	30	contacting JTAC.....	xxi
standby state, signaled LSPs.....	177, 362	tie-breaking rules, path calculation.....	18, 149, 351
standby statement.....	362	timer (tracing flag)	
usage guidelines.....	177	MPLS.....	372
state (tracing flag)		to statement	
MPLS.....	372	MPLS.....	370
static (tracing flag).....	372	usage guidelines.....	139
static LSPs		static LSPs	
configuring.....	189	usage guidelines (ingress router).....	189
egress routers.....	192, 194, 320	traceoptions statement	
ingress routers.....	189	MPLS.....	371
intermediate routers.....	192, 320	usage guidelines.....	249
overview.....	15	tracing flags	
revert timer.....	142, 193	all.....	371
static LSP metric.....	148, 329	connection.....	371
static MPLS.....	189	connection-detail.....	371
static routes		cspf.....	371
point-to-multipoint LSPs.....	195	cspf-link.....	371
static-label-switched path statement		cspf-node.....	371
usage guidelines.....	189	error	
static-label-switched-path statement		MPLS.....	371
static LSP.....	363	graceful-restart.....	372
statistics		lspring.....	372
MPLS traffic.....	233, 365	nsr-synchronization.....	372
output file.....	233	nsr-synchronization-detail.....	372
statistics statement.....	365	state	
usage guidelines.....	233	MPLS.....	372
strict explicit routes.....	6, 196	static.....	372
subscription statement		timer	
usage guidelines.....	179	MPLS.....	372
summary LSA.....	230	tracing operations	
support, technical See technical support		MPLS.....	249, 371
swap (label operation).....	13	traffic	
swap and push (label operation).....	13	protection, MPLS.....	30
		statistics.....	233, 365



traffic engineering	
BGP destinations.....	26
fate-sharing.....	19
IGP and BGP destinations.....	28
IGP shortcuts.....	21
information distribution component.....	5
inter-area, OSPF.....	230
LSP metric advertisement.....	230
overview.....	3, 10
packet-forwarding component.....	4
path-selection component.....	6
routing options.....	7
signaling component.....	7
srlg.....	20
traffic engineering database accuracy.....	57
traffic engineering database.....	15
accuracy.....	57
traffic-engineered LSPs	
fast reroute.....	183
traffic-engineering statement	
bgp-igp option.....	228
bgp-igp-both-ribs option.....	228
MPLS.....	373
usage guidelines.....	227
mpls-forwarding option.....	229
usage guidelines.....	201
transit routers.....	14
transit-lsp-association	
usage guidelines.....	123
transit-lsp-association statement	
MPLS-TP.....	374
transitions	
advertising.....	178, 292
damping.....	178
traps, SNMP See SNMP traps	
TTL decrementing	
disabling.....	153, 334, 336
tunneling, MPLS	
RSVP LSPs.....	323

## U

ultimate-hop popping.....	227
unstable LSPs	
fate-sharing See fate-sharing	

## V

verification	
network interfaces.....	218

## W

world-readable option to statistics statement	
MPLS.....	365
world-readable option to traceoptions statement	
MPLS.....	372

