

# MPLS Feature Guide for QFX10000 Switches

Release

15.1X53



Modified: 2016-10-25

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

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

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

*MPLS Feature Guide for QFX10000 Switches*  
15.1X53  
Copyright © 2016, 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 . . . . .	xix
	Documentation and Release Notes . . . . .	xix
	Supported Platforms . . . . .	xix
	Using the Examples in This Manual . . . . .	xix
	Merging a Full Example . . . . .	xx
	Merging a Snippet . . . . .	xx
	Documentation Conventions . . . . .	xxi
	Documentation Feedback . . . . .	xxiii
	Requesting Technical Support . . . . .	xxiii
	Self-Help Online Tools and Resources . . . . .	xxiii
	Opening a Case with JTAC . . . . .	xxiv
<b>Part 1</b>	<b>LDP</b>	
<b>Chapter 1</b>	<b>Using LDP . . . . .</b>	<b>3</b>
	LDP Introduction . . . . .	4
	Junos OS LDP Protocol Implementation . . . . .	4
	LDP Operation . . . . .	4
	Tunneling LDP LSPs in RSVP LSPs . . . . .	5
	Tunneling LDP LSPs in RSVP LSPs Overview . . . . .	5
	Label Operations . . . . .	6
	LDP Message Types . . . . .	7
	Discovery Messages . . . . .	7
	Session Messages . . . . .	8
	Advertisement Messages . . . . .	8
	Notification Messages . . . . .	8
	LDP Session Protection . . . . .	8
	LDP Graceful Restart . . . . .	9
	Minimum LDP Configuration . . . . .	10
	Enabling and Disabling LDP . . . . .	10
	Enabling Strict Targeted Hello Messages for LDP . . . . .	10
	Filtering Inbound LDP Label Bindings . . . . .	11
	Examples: Filtering Inbound LDP Label Bindings . . . . .	12
	Filtering Outbound LDP Label Bindings . . . . .	13
	Examples: Filtering Outbound LDP Label Bindings . . . . .	14
	Specifying the Transport Address Used by LDP . . . . .	15
	Collecting LDP Statistics . . . . .	15
	LDP Statistics Output . . . . .	16
	Disabling LDP Statistics on the Penultimate-Hop Router . . . . .	16
	LDP Statistics Limitations . . . . .	17

Tracing LDP Protocol Traffic . . . . .	17
Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels . . . . .	18
Tracing LDP Protocol Traffic Within FECs . . . . .	19
Examples: Tracing LDP Protocol Traffic . . . . .	19
Example: Configuring LDP Downstream on Demand . . . . .	20
Configuring the LDP Timer for Hello Messages . . . . .	25
Configuring the LDP Timer for Link Hello Messages . . . . .	26
Configuring the LDP Timer for Targeted Hello Messages . . . . .	26
Configuring the Delay Before LDP Neighbors Are Considered Down . . . . .	26
Configuring the LDP Hold Time for Link Hello Messages . . . . .	27
Configuring the LDP Hold Time for Targeted Hello Messages . . . . .	27
Configuring the Interval for LDP Keepalive Messages . . . . .	27
Configuring the LDP Keepalive Timeout . . . . .	28
Configuring LDP Route Preferences . . . . .	28
Configuring LDP Graceful Restart . . . . .	28
Enabling Graceful Restart . . . . .	29
Disabling LDP Graceful Restart or Helper Mode . . . . .	29
Configuring Reconnect Time . . . . .	30
Configuring Recovery Time and Maximum Recovery Time . . . . .	30
Configuring the Prefixes Advertised into LDP from the Routing Table . . . . .	31
Example: Configuring the Prefixes Advertised into LDP . . . . .	31
Configuring LDP LSP Traceroute . . . . .	32
Configuring Miscellaneous LDP Properties . . . . .	33
Configuring LDP to Use the IGP Route Metric . . . . .	33
Preventing Addition of Ingress Routes to the inet.0 Routing Table . . . . .	33
Multiple-Instance LDP and Carrier-of-Carriers VPNs . . . . .	34
Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router . . . . .	34
Enabling LDP over RSVP-Established LSPs . . . . .	34
Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks . . . . .	35
Configuring the TCP MD5 Signature for LDP Sessions . . . . .	35
Configuring LDP Session Protection . . . . .	36
Disabling SNMP Traps for LDP . . . . .	37
Configuring LDP Synchronization with the IGP on LDP Links . . . . .	37
Configuring LDP Synchronization with the IGP on the Router . . . . .	38
Configuring the Label Withdrawal Timer . . . . .	38
Ignoring the LDP Subnet Check . . . . .	38
<b>Chapter 2 Configuration Statements for LDP . . . . .</b>	<b>41</b>
allow-subnet-mismatch . . . . .	42
authentication-algorithm . . . . .	43
authentication-key (Protocols LDP) . . . . .	45
authentication-key-chain (Protocols LDP) . . . . .	46
deaggregate . . . . .	47
disable (Protocols LDP) . . . . .	48
dod-request-policy . . . . .	49
downstream-on-demand . . . . .	49
ecmp . . . . .	50

egress-policy	50
explicit-null (Protocols LDP)	51
export (Protocols LDP)	51
fec	52
graceful-restart (Protocols LDP)	53
hello-interval (Protocols LDP)	54
helper-disable (LDP)	55
hold-time (Protocols LDP)	56
ignore-lsp-metrics	57
igp-synchronization	57
import (Protocols LDP)	58
interface (Protocols LDP)	59
keepalive-interval	60
keepalive-timeout	61
l2-smart-policy	61
label-withdrawal-delay	62
ldp	63
ldp-synchronization	66
ldp-tunneling	66
log-updown (Protocols LDP)	67
maximum-neighbor-recovery-time	68
no-forwarding	69
policing (Protocols LDP)	70
preference (Protocols LDP)	71
reconnect-time	72
recovery-time	73
session (ldp)	74
session-protection	75
strict-targeted-hellos	75
targeted-hello	76
traceoptions (Protocols LDP)	77
track-igp-metric	79
traffic-statistics (Protocols LDP)	80
transport-address	82
<b>Chapter 3</b>	
<b>Monitoring Commands for LDP</b>	<b>83</b>
clear ldp neighbor	84
clear ldp session	85
clear ldp statistics	86
ping mpls ldp	87
show ldp database	90
show ldp fec-filters	99
show ldp interface	100
show ldp neighbor	102
show ldp path	104
show ldp route	106
show ldp session	110
show ldp statistics	116
show ldp traffic-statistics	120

## Part 2

### Chapter 4

traceroute mpls ldp . . . . .	124
<b>MPLS</b>	
<b>Using MPLS . . . . .</b>	<b>129</b>
MPLS Overview For QFX Series and EX4600 Switches . . . . .	130
Why Use MPLS? . . . . .	131
Why Not Use MPLS? . . . . .	131
How Do I Configure MPLS? . . . . .	131
Configure the MPLS LER (Ingress) Switch and the Egress Switch . . . . .	132
Configure LSRs for MPLS . . . . .	132
What Does the MPLS Protocol Do? . . . . .	132
How Does MPLS Interface to Other Protocols? . . . . .	133
If I Have Used Cisco MPLS, What Do I Need to Know? . . . . .	133
MPLS Feature Support on QFX Series and EX4600 Switches . . . . .	134
MPLS Commands Supported by QFX Series and EX4600 Switches . . . . .	135
MPLS Features Supported by QFX Series and EX4600 Switches . . . . .	135
MPLS Limitations on QFX Series and EX4600 Switches . . . . .	142
MPLS Limitations on QFX3500 Switches . . . . .	143
MPLS Limitations on QFX5100 and EX4600 Switches . . . . .	143
MPLS Limitations on QFX5100 Virtual Chassis and Virtual Chassis Fabric . . . . .	145
MPLS Limitations on QFX10000 Switches . . . . .	145
Understanding MPLS Components for QFX Series and EX4600 Switches . . . . .	146
Provider Edge Switches . . . . .	146
MPLS Protocol and Label-Switched Paths . . . . .	146
IP Over MPLS for Customer Edge Interfaces . . . . .	146
BGP Layer 3 VPN Configuration . . . . .	147
Routing Instances for Layer 3 VPN . . . . .	147
Routing Instances for Layer 2 VPN and Layer 3 VPN . . . . .	147
Ethernet Encapsulation for Layer 2 VPN . . . . .	147
Provider Switch . . . . .	147
Components Required for All Switches in the MPLS Network . . . . .	148
Interior Gateway Protocol . . . . .	148
Traffic Engineering . . . . .	148
MPLS Protocol . . . . .	148
RSVP . . . . .	148
Family mpls . . . . .	149
Understanding MPLS Label Operations . . . . .	150
MPLS Label-Switched Paths and MPLS Labels . . . . .	150
Reserved Labels . . . . .	151
MPLS Label Operations . . . . .	151
Penultimate-Hop Popping and Ultimate-Hop Popping . . . . .	153
Understanding BGP . . . . .	154
Autonomous Systems . . . . .	154
AS Paths and Attributes . . . . .	154
External and Internal BGP . . . . .	155
Multiple Instances of BGP . . . . .	155
IPv6 Layer 3 VPNs . . . . .	156

Ethernet Pseudowire Overview .....	157
Understanding CoS MPLS EXP Classifiers and Rewrite Rules .....	158
EXP Classifiers .....	159
EXP Rewrite Rules .....	160
Schedulers .....	161
Understanding Ethernet-over-MPLS (L2 Circuit) .....	161
Ethernet-over-MPLS in Data Centers .....	161
Understanding Using MPLS-Based Layer 3 VPNs on Switches .....	162
MPLS-Based Layer 3 VPNs .....	162
Carrier-of-Carriers VPNs .....	164
Internet Service Provider as the Customer .....	165
VPN Service Provider as the Customer .....	165
Interprovider and Carrier-of-Carriers VPNs .....	165
Chained Composite Next Hops for Transit Devices for VPNs .....	166
Fast Reroute Overview .....	168
Graceful Restart and MPLS-Related Protocols .....	170
LDP .....	170
RSVP .....	171
CCC and TCC .....	171
Types of LSPs .....	171
Configuring Automatic Bandwidth Allocation for LSPs .....	172
Configuring Automatic Bandwidth Allocation on LSPs .....	173
Configuring the Automatic Bandwidth Allocation Interval .....	174
Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth .....	174
Configuring the Automatic Bandwidth Adjustment Threshold .....	175
Configuring a Limit on Bandwidth Overflow and Underflow Samples .....	175
Configuring Passive Bandwidth Utilization Monitoring .....	177
Requesting Automatic Bandwidth Allocation Adjustment .....	178
Configuring CoS Bits for an MPLS Network .....	179
Configuring Ethernet over MPLS (L2 Circuit) .....	180
Configuring the Local PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire) .....	181
Configuring the Remote PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire) .....	181
Configuring the Local PE Switch for VLAN-Based Layer 2 Circuit .....	182
Configuring the Remote PE Switch for VLAN-Based Layer 2 Circuit .....	183
Configuring a Global MPLS EXP Classifier .....	184
Configuring MPLS Firewall Filters and Policers .....	184
Configuring an MPLS Firewall Filter .....	186
Applying an MPLS Firewall Filter to an MPLS Interface .....	186
Configuring Policers for LSPs .....	187
Configuring MPLS to Gather Statistics .....	187
Configuring MPLS on Provider Edge Switches .....	188
Configuring the Ingress PE Switch .....	189
Configuring the Egress PE Switch .....	190
Configuring MPLS on Provider Switches .....	192

	Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs	193
	Configuring Static Label Switched Paths for MPLS	196
	Configuring the Ingress PE Switch	197
	Configuring the Provider and the Egress PE Switch	198
	Configuring Rewrite Rules for MPLS EXP Classifiers	199
	Example: Configuring MPLS-Based Layer 3 VPNs	200
	Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks	209
	Verifying That MPLS Is Working Correctly	218
	Verifying the Physical Layer on the Switches	218
	Verifying the Routing Protocol	218
	Verifying the Core Interfaces Being Used for the MPLS Traffic	219
	Verifying RSVP	219
	MPLS Configuration Guidelines	220
	Supported MPLS Scaling Values	221
	MPLS Stitching For Virtual Machine Connection	222
	When Would I Use Stitching?	222
	How Does MPLS Stitching Work?	222
	How Do I Configure Stitching?	223
	Which Switches Support Stitching?	223
	Q&A	223
<b>Chapter 5</b>	<b>Configuration Statements for MPLS</b>	<b>225</b>
	adaptive	228
	adjust-interval	228
	adjust-threshold	229
	adjust-threshold-overflow-limit	229
	adjust-threshold-underflow-limit	230
	admin-down	230
	advertisement-hold-time	231
	associate-backup-pe-groups	231
	auto-bandwidth (MPLS Tunnel)	232
	backup-pe-group	233
	bandwidth (Fast Reroute, Signaled, and Multiclass LSPs)	234
	bandwidth-model	235
	bypass (Static LSP)	236
	chained-composite-next-hop	237
	class-of-service (Protocols MPLS)	239
	corouted-bidirectional	240
	corouted-bidirectional-passive	240
	description (Protocols MPLS)	241
	diffserv-te	242
	disable (Protocols MPLS)	243
	exclude (for Fast Reroute)	244
	exclude-srlg	245
	explicit-null (Protocols MPLS)	246
	fast-reroute (Protocols MPLS)	247
	forwarding-table	248
	from (Protocols MPLS)	248



gpipid	249
hop-limit	250
include-all (for Fast Reroute)	251
include-any (for Fast Reroute)	251
ingress (LSP)	252
install (Protocols MPLS)	253
interface (Protocols MPLS)	254
ipv6-tunneling	255
l2circuit	256
label-switched-path (Protocols MPLS)	258
ldp-tunneling	260
link-protection (Static LSPs)	261
log-updown (Protocols MPLS)	262
lsp-attributes	263
maximum-bandwidth (Protocols MPLS)	264
metric (Protocols MPLS)	264
minimum-bandwidth	265
monitor-bandwidth	265
mtu-signaling	266
no-cspf	267
no-decrement-ttl	268
no-install-to-address	269
no-propagate-ttl	270
record	271
no-trap	272
node-link-protection (Protocols MPLS)	273
oam (Protocols MPLS)	274
optimize-aggressive	275
optimize-hold-dead-delay	276
optimize-switchover-delay	277
optimize-timer (Protocols MPLS)	278
p2mp (Protocols MPLS)	279
path (Protocols MPLS)	280
path-mtu	281
policing (Protocols MPLS)	282
policy-statement	283
pop	287
preference (Protocols MPLS)	288
primary (Protocols MPLS)	289
push	290
record	291
retry-limit	292
revert-timer	293
rsdp-error-hold-time	294
secondary (Protocols MPLS)	295
select	296
signal-bandwidth	296
smart-optimize-timer	297
standby	298

	static-label-switched-path . . . . .	299
	statistics (Protocols MPLS) . . . . .	301
	swap . . . . .	302
	switching-type . . . . .	303
	te-class-matrix . . . . .	304
	traffic-engineering (Protocols MPLS) . . . . .	305
	transit (Chained Composite Next Hops) . . . . .	306
	transit-lsp-association . . . . .	308
	te-class-matrix . . . . .	309
	to . . . . .	310
	traceoptions (Protocols MPLS) . . . . .	311
	traffic-engineering (Protocols MPLS) . . . . .	313
	transit-lsp-association . . . . .	314
<b>Chapter 6</b>	<b>Monitoring Commands for MPLS . . . . .</b>	<b>315</b>
	clear mpls lsp . . . . .	317
	monitor label-switched-path . . . . .	319
	ping mpls bgp . . . . .	322
	ping mpls l2circuit . . . . .	324
	ping mpls l3vpn . . . . .	327
	ping mpls lsp-end-point . . . . .	330
	request mpls lsp adjust-autobandwidth . . . . .	332
	show security keychain . . . . .	334
	show link-management . . . . .	337
	show link-management peer . . . . .	341
	show link-management routing . . . . .	343
	show link-management statistics . . . . .	346
	show link-management te-link . . . . .	348
	show mpls call-admission-control . . . . .	350
	show mpls cspf . . . . .	352
	show mpls diffserv-te . . . . .	354
	show route forwarding-table . . . . .	356
	show mpls interface . . . . .	364
	show link-management statistics . . . . .	366
	show link-management te-link . . . . .	368
	show mpls call-admission-control . . . . .	370
	show mpls cspf . . . . .	372
	show mpls diffserv-te . . . . .	374
	show route forwarding-table . . . . .	376
	show mpls interface . . . . .	384
	show mpls lsp . . . . .	386
	show mpls lsp autobandwidth . . . . .	404
	show mpls path . . . . .	407
	show route table . . . . .	409
	show route forwarding-table . . . . .	437
	show mpls static-lsp . . . . .	451
	show ted database . . . . .	454
	show ted link . . . . .	462
	show ted protocol . . . . .	465

<b>Part 3</b>	<b>RSVP</b>	
<b>Chapter 7</b>	<b>Using RSVP</b>	<b>469</b>
	Understanding MPLS Components for QFX Series and EX4600 Switches . . . .	469
	Provider Edge Switches . . . . .	469
	MPLS Protocol and Label-Switched Paths . . . . .	470
	IP Over MPLS for Customer Edge Interfaces . . . . .	470
	BGP Layer 3 VPN Configuration . . . . .	470
	Routing Instances for Layer 3 VPN . . . . .	470
	Routing Instances for Layer 2 VPN and Layer 3 VPN . . . . .	470
	Ethernet Encapsulation for Layer 2 VPN . . . . .	471
	Provider Switch . . . . .	471
	Components Required for All Switches in the MPLS Network . . . . .	471
	Interior Gateway Protocol . . . . .	471
	Traffic Engineering . . . . .	472
	MPLS Protocol . . . . .	472
	RSVP . . . . .	472
	Family mpls . . . . .	472
	RSVP Overview . . . . .	473
	MTU Signaling in RSVP . . . . .	474
	Tunneling LDP LSPs in RSVP LSPs . . . . .	474
	Tunneling LDP LSPs in RSVP LSPs Overview . . . . .	475
	Configuring MPLS on Provider Edge Switches . . . . .	475
	Configuring the Ingress PE Switch . . . . .	476
	Configuring the Egress PE Switch . . . . .	477
	Configuring MPLS on Provider Switches . . . . .	479
	Verifying That MPLS Is Working Correctly . . . . .	480
	Verifying the Physical Layer on the Switches . . . . .	480
	Verifying the Routing Protocol . . . . .	481
	Verifying the Core Interfaces Being Used for the MPLS Traffic . . . . .	481
	Verifying RSVP . . . . .	481
	Dynamic Bandwidth Management Using Container LSP Overview . . . . .	482
	Understanding RSVP Multipath Extensions . . . . .	482
	Junos OS RSVP Multipath Implementation . . . . .	483
	Current Traffic Engineering Challenges . . . . .	484
	Using Container LSP as a Solution . . . . .	487
	Accommodating the New Demand X . . . . .	487
	Creating New LSPs to Meet Demand X . . . . .	488
	Assigning Bandwidth to the New LSPs . . . . .	488
	Controlling the LSP Paths . . . . .	488
	Junos OS Container LSP Implementation . . . . .	489
	Container LSP Terminology . . . . .	489
	LSP Splitting . . . . .	490
	LSP Merging . . . . .	492
	Node and Link Protection . . . . .	494
	Naming Convention . . . . .	494
	Normalization . . . . .	495
	Constraint-Based Routing Path Computation . . . . .	500
	Sampling . . . . .	501

	Support for NSR, IPG-FA, and Static Routes . . . . .	501
	Configuration Statements Supported for Container LSPs . . . . .	504
	Impact of Configuring Container LSPs on Network Performance . . . . .	508
	Supported and Unsupported Features . . . . .	509
<b>Chapter 8</b>	<b>Configuration Statements for RSVP . . . . .</b>	<b>511</b>
	admin-group . . . . .	513
	authentication-key (Protocols RSVP) . . . . .	514
	aggregate (Protocols RSVP) . . . . .	515
	bandwidth (Protocols RSVP) . . . . .	516
	bypass (Signaled LSP) . . . . .	517
	class-of-service (Protocols RSVP) . . . . .	518
	container-label-switched-path . . . . .	519
	disable (Protocols RSVP) . . . . .	520
	fast-reroute (Protocols RSVP) . . . . .	521
	graceful-deletion-timeout . . . . .	521
	graceful-restart (Enabling Globally) . . . . .	522
	hello-acknowledgements . . . . .	523
	hello-interval (Protocols RSVP) . . . . .	524
	helper-disable (Multiple Protocols) . . . . .	525
	hop-limit . . . . .	526
	interface (Protocols RSVP) . . . . .	527
	keep-multiplier . . . . .	528
	link-protection (RSVP) . . . . .	529
	load-balance (Protocols RSVP) . . . . .	530
	max-bypasses . . . . .	531
	maximum-helper-recovery-time . . . . .	532
	maximum-helper-restart-time (RSVP) . . . . .	533
	no-cspf (Protocols RSVP) . . . . .	534
	no-interface-hello . . . . .	534
	no-local-reversion . . . . .	535
	no-node-id-subobject . . . . .	536
	no-p2mp-sublsp . . . . .	536
	node-hello . . . . .	537
	optimize-timer (Protocols RSVP) . . . . .	537
	path (Protocols RSVP) . . . . .	538
	preemption . . . . .	539
	priority (Protocols RSVP) . . . . .	540
	refresh-time . . . . .	541
	reliable . . . . .	541
	setup-protection . . . . .	542
	subscription . . . . .	543
	soft-preemption (Protocols RSVP) . . . . .	544
	splitting-merging . . . . .	545
	traceoptions (Protocols RSVP) . . . . .	547
	tunnel-services (RSVP) . . . . .	549
	update-threshold . . . . .	549

<b>Chapter 9</b>	<b>Monitoring Commands for RSVP .....</b>	<b>551</b>
	clear mpls container-lsp .....	552
	clear rsvp session .....	554
	clear rsvp statistics .....	556
	ping mpls rsvp .....	557
	request mpls container-lsp .....	562
	clear mpls container-lsp .....	563
	show rsvp interface .....	565
	show rsvp neighbor .....	570
	show rsvp session .....	575
	show rsvp statistics .....	585
	show rsvp version .....	589
	traceroute mpls rsvp .....	592



# List of Figures

<b>Part 1</b>	<b>LDP</b>	
<b>Chapter 1</b>	<b>Using LDP</b> .....	<b>3</b>
	Figure 1: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs . . . .	6
	Figure 2: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs . . . . .	7
<b>Part 2</b>	<b>MPLS</b>	
<b>Chapter 4</b>	<b>Using MPLS</b> .....	<b>129</b>
	Figure 3: Label Encoding . . . . .	151
	Figure 4: MPLS Label Swapping . . . . .	152
	Figure 5: ASs, EBGp, and IBGP . . . . .	155
	Figure 6: Carrier-of-Carriers VPN Architecture . . . . .	164
	Figure 7: Detours Established for an LSP Using Fast Reroute . . . . .	168
	Figure 8: Detour After the Link from Router B to Router C Fails . . . . .	168
	Figure 9: Detours Merging into Other Detours . . . . .	169
	Figure 10: Ethernet over MPLS Layer 2 Circuit . . . . .	180
	Figure 11: Configuring an MPLS-Based Layer 3 VPN . . . . .	201
	Figure 12: IPv6 Networks Linked by MPLS IPv4 Tunnels . . . . .	210
	Figure 13: Virtual Machines on Either Side of Routers . . . . .	222
<b>Part 3</b>	<b>RSVP</b>	
<b>Chapter 7</b>	<b>Using RSVP</b> .....	<b>469</b>
	Figure 14: RSVP Reservation Request and Data Flow . . . . .	473
	Figure 15: Sample Topology . . . . .	484





# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xix</b>
	Table 1: Notice Icons . . . . .	xxi
	Table 2: Text and Syntax Conventions . . . . .	xxi
<b>Part 1</b>	<b>LDP</b>	
<b>Chapter 1</b>	<b>Using LDP</b> . . . . .	<b>3</b>
	Table 3: from Operators That Apply to LDP Received-Label Filtering . . . . .	11
	Table 4: to Operators for LDP Outbound-Label Filtering . . . . .	13
<b>Chapter 3</b>	<b>Monitoring Commands for LDP</b> . . . . .	<b>83</b>
	Table 5: show ldp database Output Fields . . . . .	91
	Table 6: show ldp fec-filters Output Fields . . . . .	99
	Table 7: show ldp interface Output Fields . . . . .	100
	Table 8: show ldp neighbor Output Fields . . . . .	102
	Table 9: show ldp path Output Fields . . . . .	104
	Table 10: show ldp route Output Fields . . . . .	106
	Table 11: show ldp session Output Fields . . . . .	110
	Table 12: show ldp statistics Output Fields . . . . .	116
	Table 13: show ldp traffic-statistics Output Fields . . . . .	121
	Table 14: traceroute mpls ldp Output Fields . . . . .	125
<b>Part 2</b>	<b>MPLS</b>	
<b>Chapter 4</b>	<b>Using MPLS</b> . . . . .	<b>129</b>
	Table 15: QFX10000 Switch MPLS Features with Junos OS Release Support . . .	135
	Table 16: QFX3500, QFX5100, and QFX5200 MPLS Features with Junos OS Release Support . . . . .	137
	Table 17: EX4600 Switch MPLS Features with Junos OS Release Support . . . .	140
	Table 18: Comparison of Interprovider and Carrier-of-Carriers VPNs . . . . .	165
	Table 19: Supported Match Conditions for MPLS Firewall Filters . . . . .	185
	Table 20: Supported Actions for MPLS Firewall Filters . . . . .	185
	Table 21: Local CE Switch in the MPLS-Based Layer 3 VPN Topology . . . . .	202
	Table 22: Remote CE Switch in the MPLS-Based Layer 3 VPN Topology . . . . .	202
	Table 23: Layer 3 VPN Components of the Local PE Switch . . . . .	202
	Table 24: Layer 3 VPN Components of the Remote PE Switch . . . . .	203
	Table 25: MPLS Scaling Values . . . . .	221
<b>Chapter 6</b>	<b>Monitoring Commands for MPLS</b> . . . . .	<b>315</b>
	Table 26: Output Control Keys for the monitor label-switched-path Command . . . . .	319
	Table 27: monitor label-switched-path Output Fields . . . . .	320

Table 28: show security keychain Output Fields . . . . .	334
Table 29: show link-management Output Fields . . . . .	337
Table 30: show link-management peer Output Fields . . . . .	341
Table 31: show link-management routing Output Fields . . . . .	343
Table 32: show link-management statistics Output Fields . . . . .	346
Table 33: show link-management te-link Output Fields . . . . .	348
Table 34: show mpls call-admission-control Output Fields . . . . .	350
Table 35: show mpls cspf Output Fields . . . . .	352
Table 36: show mpls diffserv-te Output Fields . . . . .	354
Table 37: show route forwarding-table Output Fields . . . . .	357
Table 38: show mpls interface Output Fields . . . . .	364
Table 39: show link-management statistics Output Fields . . . . .	366
Table 40: show link-management te-link Output Fields . . . . .	368
Table 41: show mpls call-admission-control Output Fields . . . . .	370
Table 42: show mpls cspf Output Fields . . . . .	372
Table 43: show mpls diffserv-te Output Fields . . . . .	374
Table 44: show route forwarding-table Output Fields . . . . .	377
Table 45: show mpls interface Output Fields . . . . .	384
Table 46: show mpls lsp Output Fields . . . . .	388
Table 47: show mpls lsp autobandwidth Output Fields . . . . .	404
Table 48: show mpls path Output Fields . . . . .	407
Table 49: show route table Output Fields . . . . .	410
Table 50: Next-hop Types Output Field Values . . . . .	415
Table 51: State Output Field Values . . . . .	417
Table 52: Communities Output Field Values . . . . .	419
Table 53: show route forwarding-table Output Fields . . . . .	440
Table 54: show mpls static-lsp Output Fields . . . . .	452
Table 55: show ted database Output Fields . . . . .	454
Table 56: show ted link Output Fields . . . . .	462
Table 57: show ted protocol Output Fields . . . . .	465

## Part 3

### Chapter 7

## RSVP

### Using RSVP . . . . . 469

Table 58: LSP Sequence Order for Bin Packing . . . . .	485
Table 59: LSP Sequence Order for Deadlock . . . . .	485
Table 60: LSP Sequence Order for Predictability . . . . .	487
Table 61: LSP Sequence Order for Predictability . . . . .	487
Table 62: Normalization with Per-LSP Autobandwidth Adjustment Changes . .	496
Table 63: Normalization with Traffic Growth . . . . .	498
Table 64: Applicability of RSVP LSPs Configuration to a Container LSP . . . . .	504

### Chapter 9

### Monitoring Commands for RSVP . . . . . 551

Table 65: show rsvp interface Output Fields . . . . .	566
Table 66: show rsvp neighbor Output Fields . . . . .	570
Table 67: show rsvp session Output Fields . . . . .	577
Table 68: show rsvp statistics Output Fields . . . . .	585
Table 69: show rsvp version Output Fields . . . . .	589
Table 70: traceroute mpls rsvp Output Fields . . . . .	593

# About the Documentation

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

## Documentation and Release Notes

---

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

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

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

## Supported Platforms

---

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

- [QFX Series](#)

## Using the Examples in This Manual

---

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

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

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

## Merging a Full Example

To merge a full example, follow these steps:

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

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

```
system {
  scripts {
    commit {
      file ex-script.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 [CLI Explorer](#).

## Documentation Conventions

[Table 1 on page xxi](#) defines notice icons used in this guide.

Table 1: Notice Icons

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

[Table 2 on page xxi](#) 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>

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

Convention	Description	Examples
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 or emphasizes important new terms.</li><li>Identifies guide 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 CLI User Guide</i></li><li>RFC 1997, <i>BGP Communities Attribute</i></li></ul>
<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; 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)	Encloses optional keywords or variables.	<b>stub &lt;default-metric <i>metric</i>&gt;;</b>
(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>  <b>(<i>string1</i>   <i>string2</i>   <i>string3</i>)</b>
# (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)	Encloses 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 ( { } )	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents 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>

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

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu 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 provide feedback by using either of the following methods:

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

## Requesting Technical Support

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

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

## Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>

- Download the latest versions of software and review release notes:  
<http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications:  
<http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum:  
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

## Opening a Case with JTAC

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

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

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



## PART 1

# LDP

- [Using LDP on page 3](#)
- [Configuration Statements for LDP on page 41](#)
- [Monitoring Commands for LDP on page 83](#)



## CHAPTER 1

# Using LDP

- [LDP Introduction on page 4](#)
- [Junos OS LDP Protocol Implementation on page 4](#)
- [LDP Operation on page 4](#)
- [Tunneling LDP LSPs in RSVP LSPs on page 5](#)
- [Tunneling LDP LSPs in RSVP LSPs Overview on page 5](#)
- [Label Operations on page 6](#)
- [LDP Message Types on page 7](#)
- [Discovery Messages on page 7](#)
- [Session Messages on page 8](#)
- [Advertisement Messages on page 8](#)
- [Notification Messages on page 8](#)
- [LDP Session Protection on page 8](#)
- [LDP Graceful Restart on page 9](#)
- [Minimum LDP Configuration on page 10](#)
- [Enabling and Disabling LDP on page 10](#)
- [Enabling Strict Targeted Hello Messages for LDP on page 10](#)
- [Filtering Inbound LDP Label Bindings on page 11](#)
- [Filtering Outbound LDP Label Bindings on page 13](#)
- [Specifying the Transport Address Used by LDP on page 15](#)
- [Collecting LDP Statistics on page 15](#)
- [Tracing LDP Protocol Traffic on page 17](#)
- [Example: Configuring LDP Downstream on Demand on page 20](#)
- [Configuring the LDP Timer for Hello Messages on page 25](#)
- [Configuring the Delay Before LDP Neighbors Are Considered Down on page 26](#)
- [Configuring the Interval for LDP Keepalive Messages on page 27](#)
- [Configuring the LDP Keepalive Timeout on page 28](#)
- [Configuring LDP Route Preferences on page 28](#)
- [Configuring LDP Graceful Restart on page 28](#)

- [Configuring the Prefixes Advertised into LDP from the Routing Table on page 31](#)
- [Configuring LDP LSP Traceroute on page 32](#)
- [Configuring Miscellaneous LDP Properties on page 33](#)

## LDP Introduction

---

The Label Distribution Protocol (LDP) is a protocol for distributing labels in non-traffic-engineered applications. LDP allows routers to establish label-switched paths (LSPs) through a network by mapping network-layer routing information directly to data link layer-switched paths.

These LSPs might have an endpoint at a directly attached neighbor (comparable to IP hop-by-hop forwarding), or at a network egress node, enabling switching through all intermediary nodes. LSPs established by LDP can also traverse traffic-engineered LSPs created by RSVP.

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers. This process forms a tree of LSPs that converge on the egress router.

## Junos OS LDP Protocol Implementation

---

The Junos OS implementation of LDP supports LDP version 1. The Junos OS supports a simple mechanism for tunneling between routers in an interior gateway protocol (IGP), to eliminate the required distribution of external routes within the core. The Junos OS allows an MPLS tunnel next hop to all egress routers in the network, with only an IGP running in the core to distribute routes to egress routers. Edge routers run BGP but do not distribute external routes to the core. Instead, the recursive route lookup at the edge resolves to an LSP switched to the egress router. No external routes are necessary on the transit LDP routers.

## LDP Operation

---

You must configure LDP for each interface on which you want LDP to run. LDP creates LSP trees rooted at each egress router for the router ID address that is the subsequent BGP next hop. The ingress point is at every router running LDP. This process provides an inet.3 route to every egress router. If BGP is running, it will attempt to resolve next hops by using the inet.3 table first, which binds most, if not all, of the BGP routes to MPLS tunnel next hops.

Two adjacent routers running LDP become neighbors. If the two routers are connected by more than one interface, they become neighbors on each interface. When LDP routers become neighbors, they establish an LDP session to exchange label information. If per-router labels are in use on both routers, only one LDP session is established between them, even if they are neighbors on multiple interfaces. For this reason, an LDP session is not related to a particular interface.

LDP operates in conjunction with a unicast routing protocol. LDP installs LSPs only when both LDP and the routing protocol are enabled. For this reason, you must enable both LDP and the routing protocol on the same set of interfaces. If this is not done, LSPs might not be established between each egress router and all ingress routers, which might result in loss of BGP-routed traffic.

You can apply policy filters to labels received from and distributed to other routers through LDP. Policy filters provide you with a mechanism to control the establishment of LSPs.

For LDP to run on an interface, MPLS must be enabled on a logical interface on that interface. For more information, see the *Logical Interfaces*.

**Related Documentation**

- [Logical Interfaces](#)

## Tunneling LDP LSPs in RSVP LSPs

You can tunnel LDP LSPs over RSVP LSPs. The following sections describe how tunneling of LDP LSPs in RSVP LSPs works:

- [Tunneling LDP LSPs in RSVP LSPs Overview on page 5](#)
- [Label Operations on page 6](#)

## Tunneling LDP LSPs in RSVP LSPs Overview

If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.

When you configure the router to run LDP across RSVP-established LSPs, LDP automatically establishes sessions with the router at the other end of the LSP. LDP control packets are routed hop-by-hop, rather than carried through the LSP. This routing allows you to use simplex (one-way) traffic-engineered LSPs. Traffic in the opposite direction flows through LDP-established LSPs that follow unicast routing rather than through traffic-engineered tunnels.

If you configure LDP over RSVP LSPs, you can still configure multiple OSPF areas and IS-IS levels in the traffic engineered core and in the surrounding LDP cloud.



**NOTE:** Beginning with Junos OS Release 15.1, multi-instance support is extended to LDP over RSVP tunneling for a virtual router routing instance. This allows splitting of a single routing and MPLS domain into multiple domains so that each domain can be scaled independently. BGP labeled unicast can be used to stitch these domains for service FECs. Each domain uses intra-domain LDP over RSVP LSP for MPLS forwarding.

- Related Documentation**
- [Label Operations on page 6](#)
  - [Configuring a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs Over a Single RSVP LSP](#)

## Label Operations

Figure 1 on page 6 depicts an LDP LSP being tunneled through an RSVP LSP. (For definitions of label operations, see *MPLS Label Overview*.) The shaded inner oval represents the RSVP domain, whereas the outer oval depicts the LDP domain. RSVP establishes an LSP through routers B, C, D, and E, with the sequence of labels L3, L4. LDP establishes an LSP through Routers A, B, E, F, and G, with the sequence of labels L1, L2, L5. LDP views the RSVP LSP between Routers B and E as a single hop.

When the packet arrives at Router A, it enters the LSP established by LDP, and a label (L1) is pushed onto the packet. When the packet arrives at Router B, the label (L1) is swapped with another label (L2). Because the packet is entering the traffic-engineered LSP established by RSVP, a second label (L3) is pushed onto the packet.

This outer label (L3) is swapped with a new label (L4) at the intermediate router (C) within the RSVP LSP tunnel, and when the penultimate router (D) is reached, the top label is popped. Router E swaps the label (L2) with a new label (L5), and the penultimate router for the LDP-established LSP (F) pops the last label.

Figure 1: Swap and Push When LDP LSPs Are Tunneled Through RSVP LSPs





Figure 2 on page 7 depicts a double push label operation (L1L2). A double push label operation is used when the ingress router (A) for both the LDP LSP and the RSVP LSP tunneled through it is the same device. Note that Router D is the penultimate hop for the LDP-established LSP, so L2 is popped from the packet by Router D.

Figure 2: Double Push When LDP LSPs Are Tunneled Through RSVP LSPs



## LDP Message Types

LDP uses the message types described in the following sections to establish and remove mappings and to report errors. All LDP messages have a common structure that uses a type, length, and value (TLV) encoding scheme.

- [Discovery Messages on page 7](#)
- [Session Messages on page 8](#)
- [Advertisement Messages on page 8](#)
- [Notification Messages on page 8](#)

## Discovery Messages

Discovery messages announce and maintain the presence of a router in a network. Routers indicate their presence in a network by sending hello messages periodically. Hello messages are transmitted as UDP packets to the LDP port at the group multicast address for all routers on the subnet.

LDP uses the following discovery procedures:

- **Basic discovery**—A router periodically sends LDP link hello messages through an interface. LDP link hello messages are sent as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.
- **Extended discovery**—LDP sessions between routers not directly connected are supported by LDP extended discovery. A router periodically sends LDP targeted hello messages to a specific address. Targeted hello messages are sent as UDP packets addressed to the LDP discovery port at the specific address. The targeted router decides whether to respond to or ignore the targeted hello message. A targeted router that chooses to respond does so by periodically sending targeted hello messages to the initiating router.

## Session Messages

---

Session messages establish, maintain, and terminate sessions between LDP peers. When a router establishes a session with another router learned through the hello message, it uses the LDP initialization procedure over TCP transport. When the initialization procedure is completed successfully, the two routers are LDP peers and can exchange advertisement messages.

## Advertisement Messages

---

Advertisement messages create, change, and delete label mappings for forwarding equivalence classes (FECs). Requesting a label or advertising a label mapping to a peer is a decision made by the local router. In general, the router requests a label mapping from a neighboring router when it needs one and advertises a label mapping to a neighboring router when it wants the neighbor to use a label.

## Notification Messages

---

Notification messages provide advisory information and signal error information. LDP sends notification messages to report errors and other events of interest. There are two kinds of LDP notification messages:

- Error notifications, which signal fatal errors. If a router receives an error notification from a peer for an LDP session, it terminates the LDP session by closing the TCP transport connection for the session and discarding all label mappings learned through the session.
- Advisory notifications, which pass information to a router about the LDP session or the status of some previous message received from the peer.

## LDP Session Protection

---

LDP session protection is based on the LDP targeted hello functionality defined in RFC 5036, *LDP Specification*, and is supported by the Junos OS as well as the LDP implementations of most other vendors. It involves sending unicast User Datagram Protocol (UDP) hello packets to a remote neighbor address and receiving similar packets from the neighbor router.

If you configure LDP session protection on a router, the LDP sessions are maintained as follows:

1. An LDP session is established between a router and a remote neighboring router.
2. If all of the direct links between the routers go down, the LDP session remains up so long as there is IP connectivity between the routers based on another connection over the network.
3. When the direct link between the routers is reestablished, the LDP session is not restarted. The routers simply exchange LDP hellos with each other over the direct link.



They can then begin forwarding LDP-signaled MPLS packets using the original LDP session.

By default, LDP targeted hellos are set to the remote neighbor so long as the LDP session is up, even if there are no more link neighbors to that router. You can also specify the duration you would like to maintain the remote neighbor connection in the absence of link neighbors. When the last link neighbor for a session goes down, the Junos OS starts an LDP session protection timer. If this timer expires before any of the link neighbors come back up, the remote neighbor connection is taken down and the LDP session is terminated. If you configure a different value for the timer while it is currently running, the Junos OS updates the timer to the specified value without disrupting the current state of the LDP session.

## LDP Graceful Restart

---

LDP graceful restart enables a router whose LDP control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. It also enables a router on which helper mode is enabled to assist a neighboring router that is attempting to restart LDP.

During session initialization, a router advertises its ability to perform LDP graceful restart or to take advantage of a neighbor performing LDP graceful restart by sending the graceful restart TLV. This TLV contains two fields relevant to LDP graceful restart: the reconnect time and the recovery time. The values of the reconnect and recovery times indicate the graceful restart capabilities supported by the router.

When a router discovers that a neighboring router is restarting, it waits until the end of the recovery time before attempting to reconnect. The recovery time is the length of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This time period is also typically the length of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

You can configure LDP graceful restart in both the master instance for the LDP protocol and for a specific routing instance. You can disable graceful restart at the global level for all protocols, at the protocol level for LDP only, and on a specific routing instance. LDP graceful restart is disabled by default, because at the global level, graceful restart is disabled by default. However, helper mode (the ability to assist a neighboring router attempting a graceful restart) is enabled by default.

The following are some of the behaviors associated with LDP graceful restart:

- Outgoing labels are not maintained in restarts. New outgoing labels are allocated.
- When a router is restarting, no label-map messages are sent to neighbors that support graceful restart until the restarting router has stabilized (label-map messages are immediately sent to neighbors that do not support graceful restart). However, all other messages (keepalive, address-message, notification, and release) are sent as usual. Distributing these other messages prevents the router from distributing incomplete information.

- Helper mode and graceful restart are independent. You can disable graceful restart in the configuration, but still allow the router to cooperate with a neighbor attempting to restart gracefully.

## Minimum LDP Configuration

---

To enable LDP on a single interface, include the **ldp** statement and specify the interface using the **interface** statement. This is the minimum LDP configuration. All other LDP configuration statements are optional.

```
ldp {  
  interface interface-name;  
}
```

To enable LDP on all interfaces, specify **all** for *interface-name*.

For a list of hierarchy levels at which you can include these statements, see the statement summary sections.

## Enabling and Disabling LDP

---

LDP is routing-instance-aware. To enable LDP on a specific interface, include the following statements:

```
ldp {  
  interface interface-name;  
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections.

To enable LDP on all interfaces, specify **all** for *interface-name*.

If you have configured interface properties on a group of interfaces and want to disable LDP on one of the interfaces, include the **interface** statement with the **disable** option:

```
interface interface-name {  
  disable;  
}
```

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

## Enabling Strict Targeted Hello Messages for LDP

---

Use strict targeted hello messages to prevent LDP sessions from being established with remote neighbors that have not been specifically configured. If you configure the **strict-targeted-hellos** statement, an LDP peer does not respond to targeted hello messages coming from a source that is not one of its configured remote neighbors. Configured remote neighbors can include:

- Endpoints of RSVP tunnels for which LDP tunneling is configured
- Layer 2 circuit neighbors

If an unconfigured neighbor sends a hello message, the LDP peer ignores the message and logs an error (with the **error** trace flag) indicating the source. For example, if the LDP peer received a targeted hello from the Internet address 10.0.0.1 and no neighbor with this address is specifically configured, the following message is printed to the LDP log file:

```
LDP: Ignoring targeted hello from 10.0.0.1
```

To enable strict targeted hello messages, include the **strict-targeted-hellos** statement:

```
strict-targeted-hellos;
```

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

## Filtering Inbound LDP Label Bindings

You can filter received LDP label bindings, applying policies to accept or deny bindings advertised by neighboring routers. To configure received-label filtering, include the **import** statement:

```
import [ policy-names ];
```

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

The named policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings received from all LDP neighbors. All filtering is done with **from** statements. [Table 3 on page 11](#) lists the only **from** operators that apply to LDP received-label filtering.

**Table 3: from Operators That Apply to LDP Received-Label Filtering**

from Operator	Description
<b>interface</b>	Matches on bindings received from a neighbor that is adjacent over the specified interface
<b>neighbor</b>	Matches on bindings received from the specified LDP router ID
<b>next-hop</b>	Matches on bindings received from a neighbor advertising the specified interface address
<b>route-filter</b>	Matches on bindings with the specified prefix

If a binding is filtered, it still appears in the LDP database, but is not considered for installation as part of a label-switched path (LSP).

Generally, applying policies in LDP can be used only to block the establishment of LSPs, not to control their routing. This is because the path that an LSP follows is determined by unicast routing, and not by LDP. However, when there are multiple equal-cost paths to the destination through different neighbors, you can use LDP filtering to exclude some

of the possible next hops from consideration. (Otherwise, LDP chooses one of the possible next hops at random.)

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router (not per-interface) labels; so if multiple parallel links exist between two routers, only one LDP session is established, and it is not bound to a single interface. When a router has multiple adjacencies to the same neighbor, take care to ensure that the filter does what is expected. (Generally, using **next-hop** and **interface** is not appropriate in this case.)

If a label has been filtered (meaning that it has been rejected by the policy and is not used to construct an LSP), it is marked as filtered in the database:

```
user@host> show ldp database
Input label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.6/32 (Filtered)
Output label database, 10.10.255.1:0-10.10.255.6:0
Label Prefix
3 10.10.255.1/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

## Examples: Filtering Inbound LDP Label Bindings

Accept only /32 prefixes from all neighbors:

```
[edit]
protocols {
  ldp {
    import only-32;
    ...
  }
}
policy-options {
  policy-statement only-32 {
    term first {
      from {
        route-filter 0.0.0.0/0 upto /31;
      }
      then reject;
    }
    then accept;
  }
}
```

Accept 131.108/16 or longer from router ID 10.10.255.2 and accept all prefixes from all other neighbors:

```
[edit]
protocols {
  ldp {
    import nosy-neighbor;
    ...
  }
}
```

```

}
policy-options {
  policy-statement nosy-neighbor {
    term first {
      from {
        neighbor 10.10.255.2;
        route-filter 131.108.0.0/16 orlonger accept;
        route-filter 0.0.0.0/0 orlonger reject;
      }
    }
    then accept;
  }
}

```

## Filtering Outbound LDP Label Bindings

You can configure export policies to filter LDP outbound labels. You can filter outbound label bindings by applying routing policies to block bindings from being advertised to neighboring routers. To configure outbound label filtering, include the **export** statement:

**export** [*policy-name*];

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

The named export policy (configured at the **[edit policy-options]** hierarchy level) is applied to all label bindings transmitted to all LDP neighbors. The only **from** operator that applies to LDP outbound label filtering is **route-filter**, which matches bindings with the specified prefix. The only **to** operators that apply to outbound label filtering are the operators in [Table 4 on page 13](#).

**Table 4: to Operators for LDP Outbound-Label Filtering**

to Operator	Description
<b>interface</b>	Matches on bindings sent to a neighbor that is adjacent over the specified interface
<b>neighbor</b>	Matches on bindings sent to the specified LDP router ID
<b>next-hop</b>	Matches on bindings sent to a neighbor advertising the specified interface address

If a binding is filtered, the binding is not advertised to the neighboring router, but it can be installed as part of an LSP on the local router. You can apply policies in LDP to block the establishment of LSPs, but not to control their routing. The path an LSP follows is determined by unicast routing, not by LDP.

LDP sessions are not bound to interfaces or interface addresses. LDP advertises only per-router (not per-interface) labels. If multiple parallel links exist between two routers, only one LDP session is established, and it is not bound to a single interface.

Do not use the **next-hop** and **interface** operators when a router has multiple adjacencies to the same neighbor.

Filtered labels are marked in the database:

```
user@host> show ldp database
Input label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
100007 10.10.255.2/32
3 10.10.255.3/32
Output label database, 10.10.255.1:0-10.10.255.3:0
Label Prefix
3 10.10.255.1/32
100001 10.10.255.6/32 (Filtered)
```

For more information about how to configure policies for LDP, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

### Examples: Filtering Outbound LDP Label Bindings

Block transmission of the route for **10.10.255.6/32** to any neighbors:

```
[edit protocols]
ldp {
  export block-one;
}
policy-options {
  policy-statement block-one {
    term first {
      from {
        route-filter 10.10.255.6/32 exact;
      }
      then reject;
    }
    then accept;
  }
}
```

Send only **131.108/16** or longer to router ID **10.10.255.2**, and send all prefixes to all other routers:

```
[edit protocols]
ldp {
  export limit-lsps;
}
policy-options {
  policy-statement limit-lsps {
    term allow-one {
      from {
        route-filter 131.108.0.0/16 orlonger;
      }
      to {
        neighbor 10.10.255.2;
      }
      then accept;
    }
    term block-the-rest {
```

```
        to {  
            neighbor 10.10.255.2;  
        }  
        then reject;  
    }  
    then accept;  
}
```

## Specifying the Transport Address Used by LDP

---

Routers must first establish a TCP session between each other before they can establish an LDP session. The TCP session enables the routers to exchange the label advertisements needed for the LDP session. To establish the TCP session, each router must learn the other router's transport address. The transport address is an IP address used to identify the TCP session over which the LDP session will run.

To configure the LDP transport address, include the `transport-address` statement:

```
transport-address (router-id | interface);
```

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

If you specify the **router-id** option, the address of the router identifier is used as the transport address (unless otherwise configured, the router identifier is typically the same as the loopback address). If you specify the **interface** option, the interface address is used as the transport address for any LDP sessions to neighbors that can be reached over that interface. Note that the router identifier is used as the transport address by default.

You cannot specify the **interface** option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the **router-id** option.

**Related Documentation**

- [transport-address on page 82](#)

## Collecting LDP Statistics

---

LDP traffic statistics show the volume of traffic that has passed through a particular FEC on a router.

When you configure the **traffic-statistics** statement at the **[edit protocols ldp]** hierarchy level, the LDP traffic statistics are gathered periodically and written to a file. You can configure how often statistics are collected (in seconds) by using the **interval** option. The default collection interval is 5 minutes. You must configure an LDP statistics file; otherwise, LDP traffic statistics are not gathered. If the LSP goes down, the LDP statistics are reset.

To collect LDP traffic statistics, include the **traffic-statistics** statement:

```
traffic-statistics {
```

```

file filename <files number> <size size> <world-readable | no-world-readable>;
interval interval;
no-penultimate-hop;
}

```

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

This section includes the following topics:

- [LDP Statistics Output on page 16](#)
- [Disabling LDP Statistics on the Penultimate-Hop Router on page 16](#)
- [LDP Statistics Limitations on page 17](#)

## LDP Statistics Output

The following sample output is from an LDP statistics file:

FEC	Type	Packets	Bytes	Shared
10.255.350.448/32	Transit	0	0	No
	Ingress	0	0	No
10.255.350.450/32	Transit	0	0	Yes
	Ingress	0	0	No
10.255.350.451/32	Transit	0	0	No
	Ingress	0	0	No
220.220.220.1/32	Transit	0	0	Yes
	Ingress	0	0	No
220.220.220.2/32	Transit	0	0	Yes
	Ingress	0	0	No
220.220.220.3/32	Transit	0	0	Yes
	Ingress	0	0	No

May 28 15:02:05, read 12 statistics in 00:00:00 seconds

The LDP statistics file includes the following columns of data:

- **read**—Number of bytes of data passed by the FEC since its LSP came up.
- **read**—FEC for which LDP traffic statistics are collected.
- **read**—Number of packets passed by the FEC since its LSP came up.
- **read**—This number (which appears next to the date and time) might differ from the actual number of the statistics displayed. Some of the statistics are summarized before being displayed.
- **Shared**—A **Yes** value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such.
- **Type**—Type of traffic originating from a router, either **Ingress** (originating from this router) or **Transit** (forwarded through this router).

## Disabling LDP Statistics on the Penultimate-Hop Router

Gathering LDP traffic statistics at the penultimate-hop router can consume excessive system resources, on next-hop routes in particular. This problem is exacerbated if you have configured the **deaggregate** statement in addition to the **traffic-statistics** statement.



For routers reaching their limit of next-hop route usage, we recommend configuring the **no-penultimate-hop** option for the **traffic-statistics** statement:

```
traffic-statistics {
  no-penultimate-hop;
}
```

For a list of hierarchy levels at which you can configure the **traffic-statistics** statement, see the statement summary section for this statement.



**NOTE:** When you configure the **no-penultimate-hop** option, no statistics are available for the FECs that are the penultimate hop for this router.

Whenever you include or remove this option from the configuration, the LDP sessions are taken down and then restarted.

The following sample output is from an LDP statistics file showing routers on which the **no-penultimate-hop** option is configured:

FEC	Type	Packets	Bytes	Shared
10.255.245.218/32	Transit	0	0	No
	Ingress	4	246	No
10.255.245.221/32	Transit	statistics disabled		
	Ingress	statistics disabled		
13.1.1.0/24	Transit	statistics disabled		
	Ingress	statistics disabled		
13.1.3.0/24	Transit	statistics disabled		
	Ingress	statistics disabled		

## LDP Statistics Limitations

The following are issues related to collecting LDP statistics by configuring the **traffic-statistics** statement:

- You cannot clear the LDP statistics.
- If you shorten the specified interval, a new LDP statistics request is issued only if the statistics timer expires later than the new interval.
- A new LDP statistics collection operation cannot start until the previous one has finished. If the interval is short or if the number of LDP statistics is large, the time gap between the two statistics collections might be longer than the interval.

When an LSP goes down, the LDP statistics are reset.

## Tracing LDP Protocol Traffic

The following sections describe how to configure the trace options to examine LDP protocol traffic:

- [Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels on page 18](#)
- [Tracing LDP Protocol Traffic Within FECs on page 19](#)
- [Examples: Tracing LDP Protocol Traffic on page 19](#)

## Tracing LDP Protocol Traffic at the Protocol and Routing Instance Levels

To trace LDP protocol traffic, you can specify options in the global **traceoptions** statement at the **[edit routing-options]** hierarchy level, and you can specify LDP-specific options by including the **traceoptions** statement:

```
traceoptions {  
  file filename <files number> <size size> <world-readable | no-world-readable>;  
  flag flag <flag-modifier> <disable>;  
}
```

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

Use the **file** statement to specify the name of the file that receives the output of the tracing operation. All files are placed in the directory `/var/log`. We recommend that you place LDP-tracing output in the file **ldp-log**.

The following trace flags display the operations associated with the sending and receiving of various LDP messages. Each can carry one or more of the following modifiers:

- **address**—Trace the operation of address and address withdrawal messages.
- **binding**—Trace label-binding operations.
- **error**—Trace error conditions.
- **event**—Trace protocol events.
- **initialization**—Trace the operation of initialization messages.
- **label**—Trace the operation of label request, label map, label withdrawal, and label release messages.
- **notification**—Trace the operation of notification messages.
- **packets**—Trace the operation of address, address withdrawal, initialization, label request, label map, label withdrawal, label release, notification, and periodic messages. This modifier is equivalent to setting the **address**, **initialization**, **label**, **notification**, and **periodic** modifiers.

You can also configure the **filter** flag modifier with the **match-on address** sub-option for the **packets** flag. This allows you to trace based on the source and destination addresses of the packets.

- **path**—Trace label-switched path operations.
- **path**—Trace label-switched path operations.
- **periodic**—Trace the operation of hello and keepalive messages.
- **route**—Trace the operation of route messages.
- **state**—Trace protocol state transitions.

## Tracing LDP Protocol Traffic Within FECs

LDP associates a forwarding equivalence class (FEC) with each LSP it creates. The FEC associated with an LSP specifies which packets are mapped to that LSP. LSPs are extended through a network as each router chooses the label advertised by the next hop for the FEC and splices it to the label it advertises to all other routers.

You can trace LDP protocol traffic within a specific FEC and filter LDP trace statements based on an FEC. This is useful when you want to trace or troubleshoot LDP protocol traffic associated with an FEC. The following trace flags are available for this purpose: **route**, **path**, and **binding**.

The following example illustrates how you might configure the LDP **traceoptions** statement to filter LDP trace statements based on an FEC:

```
[edit protocols ldp traceoptions]
set flag route filter match-on fec policy "filter-policy-for-ldp-fec";
```

This feature has the following limitations:

- The filtering capability is only available for FECs composed of IP version 4 (IPv4) prefixes.
- Layer 2 circuit FECs cannot be filtered.
- When you configure both route tracing and filtering, MPLS routes are not displayed (they are blocked by the filter).
- Filtering is determined by the policy and the configured value for the **match-on** option. When configuring the policy, be sure that the default behavior is always **reject**.
- The only **match-on** option is **fec**. Consequently, the only type of policy you should include is a route-filter policy.

## Examples: Tracing LDP Protocol Traffic

Trace LDP path messages in detail:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5;
      flag path;
    }
  }
}
```

Trace all LDP outgoing messages:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5;
```

```
        flag packets;
    }
}
```

Trace all LDP error conditions:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5;
      flag error;
    }
  }
}
```

Trace all LDP incoming messages and all label-binding operations:

```
[edit]
protocols {
  ldp {
    traceoptions {
      file ldp size 10m files 5 world-readable;
      flag packets receive;
      flag binding;
    }
    interface all {
    }
  }
}
```

Trace LDP protocol traffic for an FEC associated with the LSP:

```
[edit]
protocols {
  ldp {
    traceoptions {
      flag route filter match-on fec policy filter-policy-for-ldp-fec;
    }
  }
}
```

---

## Example: Configuring LDP Downstream on Demand

This example shows how to configure LDP downstream on demand. LDP is commonly configured using downstream unsolicited advertisement mode, meaning label advertisements for all routes are received from all LDP peers. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between the access and aggregation networks and to reduce the processing requirements for the control plane.

Downstream nodes could potentially receive tens of thousands of label bindings from upstream aggregation nodes. Instead of learning and storing all label bindings for all possible loopback addresses within the entire MPLS network, the downstream aggregation node can be configured using LDP downstream on demand to only request the label

bindings for the FECs corresponding to the loopback addresses of those egress nodes on which it has services configured.

- [Requirements on page 21](#)
- [Overview on page 21](#)
- [Configuration on page 21](#)
- [Verification on page 24](#)

## Requirements

This example uses the following hardware and software components:

- M Series router
- Junos OS 12.2

## Overview

You can enable LDP downstream on demand label advertisement for an LDP session by including the [downstream-on-demand](#) statement at the [\[edit protocols ldp session\]](#) hierarchy level. If you have configured downstream on demand, the Juniper Networks router advertises the downstream on demand request to its peer routers. For a downstream on demand session to be established between two routers, both have to advertise downstream on demand mode during LDP session establishment. If one router advertises downstream unsolicited mode and the other advertises downstream on demand, downstream unsolicited mode is used.

## Configuration

### Configuring LDP Downstream on Demand

#### Step-by-Step Procedure

To configure a LDP downstream on demand policy and then configure that policy and enable LDP downstream on demand on the LDP session:

1. Configure the downstream on demand policy (DOD-Request-Loopbacks in this example).

This policy causes the router to forward label request messages only to the FECs that are matched by the DOD-Request-Loopbacks policy.

```
[edit policy-options]
user@host# set prefix-list Request-Loopbacks 10.1.1.1/32
user@host# set prefix-list Request-Loopbacks 10.1.1.2/32
user@host# set prefix-list Request-Loopbacks 10.1.1.3/32
user@host# set prefix-list Request-Loopbacks 10.1.1.4/32
user@host# set policy-statement DOD-Request-Loopbacks term 1 from prefix-list
Request-Loopbacks
user@host# set policy-statement DOD-Request-Loopbacks term 1 then accept
```

2. Specify the DOD-Request-Loopbacks policy using the [dod-request-policy](#) statement at the [\[edit protocols ldp\]](#) hierarchy level.

The policy specified with the [dod-request-policy](#) statement is used to identify the prefixes to send label request messages. This policy is similar to an egress policy

or an import policy. When processing routes from the inet.0 routing table, the Junos OS software checks for routes matching the **DOD-Request-Loopbacks** policy (in this example). If the route matches the policy and the LDP session is negotiated with DOD advertisement mode, label request messages are sent to the corresponding downstream LDP session.

```
[edit protocols ldp]
user@host# set dod-request-policy DOD-Request-Loopbacks
```

3. Include the **downstream-on-demand** statement in the configuration for the LDP session to enable downstream on demand distribution mode.

```
[edit protocols ldp]
user@host# set session 1.1.1.1 downstream-on-demand
```

### Distributing LDP Downstream on Demand Routes into Labeled BGP

**Step-by-Step Procedure** To distribute LDP downstream on demand routes into labeled BGP, use a BGP export policy.

1. Configure the LDP route policy (**redistribute\_ldp** in this example).

```
[edit policy-options]
user@host# set policy-statement redistribute_ldp term 1 from protocol ldp
user@host# set policy-statement redistribute_ldp term 1 from tag 1000
user@host# set policy-statement redistribute_ldp term 1 then accept
```

2. Include the LDP route policy, **redistribute\_ldp** in the BGP configuration (as a part of the BGP group configuration **ebgp-to-abr** in this example).

BGP forwards the LDP routes based on the **redistribute\_ldp** policy to the remote PE router

```
[edit protocols bgp]
user@host# set group ebgp-to-abr type external
user@host# set group ebgp-to-abr local-address 192.168.0.1
user@host# set group ebgp-to-abr peer-as 65319
user@host# set group ebgp-to-abr local-as 65320
user@host# set group ebgp-to-abr neighbor 192.168.6.1 family inet unicast
user@host# set group ebgp-to-abr neighbor 192.168.6.1 family inet labeled-unicast
rib inet.3
user@host# set group ebgp-to-abr neighbor 192.168.6.1 export redistribute_ldp
```

**Step-by-Step Procedure** To restrict label propagation to other routers configured in downstream unsolicited mode (instead of downstream on demand), configure the following policies:

1. Configure the **dod-routes** policy to accept routes from LDP.

```
user@host# set policy-options policy-statement dod-routes term 1 from protocol ldp
user@host# set policy-options policy-statement dod-routes term 1 from tag 1145307136
user@host# set policy-options policy-statement dod-routes term 1 then accept
```

2. Configure the **do-not-propagate-du-sessions** policy to not forward routes to neighbors 1.1.1.1, 2.2.2.2, and 3.3.3.3.

```

user@host# set policy-options policy-statement do-not-propagate-du-sessions
term 1 to neighbor 1.1.1.1
user@host# set policy-options policy-statement do-not-propagate-du-sessions
term 1 to neighbor 2.2.2.2
user@host# set policy-options policy-statement do-not-propagate-du-sessions
term 1 to neighbor 3.3.3.3
user@host# set policy-options policy-statement do-not-propagate-du-sessions
term 1 then reject

```

3. Configure the **filter-dod-on-du-sessions** policy to prevent the routes examined by the **dod-routes** policy from being forwarded to the neighboring routers defined in the **do-not-propagate-du-sessions** policy.

```

user@host# set policy-options policy-statement filter-dod-routes-on-du-sessions
term 1 from policy dod-routes
user@host# set policy-options policy-statement filter-dod-routes-on-du-sessions
term 1 to policy do-not-propagate-du-sessions

```

4. Specify the **filter-dod-routes-on-du-session** policy as the export policy for BGP group **ebgp-to-abr**.

```

[edit protocols bgp]
user@host# set group ebgp-to-abr neighbor 192.168.6.2 export
filter-dod-routes-on-du-sessions

```

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

```

user@host#
show policy-options
prefix-list Request-Loopbacks {
  10.1.1.1/32;
  10.1.1.2/32;
  10.1.1.3/32;
  10.1.1.4/32;
}
policy-statement DOD-Request-Loopbacks {
  term 1 {
    from {
      prefix-list Request-Loopbacks;
    }
    then accept;
  }
}
policy-statement redistribute_ldp {
  term 1 {
    from {
      protocol ldp;
      tag 1000;
    }
    then accept;
  }
}

user@host#
show protocols ldp

```

```
dod-request-policy DOD-Request-Loopbacks;
session 1.1.1.1 {
    downstream-on-demand;
}

user@host#
show protocols bgp
group ebgp-to-abr {
    type external;
    local-address 192.168.0.1;
    peer-as 65319;
    local-as 65320;
    neighbor 192.168.6.1 {
        family inet {
            unicast;
            labeled-unicast {
                rib {
                    inet.3;
                }
            }
        }
    }
    export redistribute_ldp;
}
```

## Verification

---

### Verifying Label Advertisement Mode

**Purpose** Confirm that the configuration is working properly.

Use the **show ldp session** command to verify the status of the label advertisement mode for the LDP session.



**Action** Issue the `show ldp session` and `show ldp session detail` commands:

- The following command output for the `show ldp session` command indicates that the **Adv. Mode** (label advertisement mode) is **DOD** (meaning the LDP downstream on demand session is operational):

```
user@host> show ldp session
  Address          State      Connection    Hold time  Adv. Mode
  1.1.1.2          Operational Open          22         DOD
```

- The following command output for the `show ldp session detail` command indicates that the **Local Label Advertisement mode** is **Downstream unsolicited**, the default value (meaning downstream on demand is not configured on the local session). Conversely, the **Remote Label Advertisement mode** and the **Negotiated Label Advertisement mode** both indicate that **Downstream on demand** is configured on the remote session

```
user@host> show ldp session detail
Address: 1.1.1.2, State: Operational, Connection: Open, Hold time: 24
Session ID: 1.1.1.1:0--1.1.1.2:0
Next keepalive in 4 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: configured-tunneled
Keepalive interval: 10, Connect retry interval: 1
Local address: 1.1.1.1, Remote address: 1.1.1.2
Up for 17:54:52
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: disabled, Helper mode: enabled,
Remote - Restart: disabled, Helper mode: enabled
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream on demand
Negotiated Label Advertisement mode: Downstream on demand
Nonstop routing state: Not in sync
Next-hop addresses received:
  1.1.1.2
```

## Configuring the LDP Timer for Hello Messages

LDP hello messages enable LDP nodes to discover one another and to detect the failure of a neighbor or the link to the neighbor. Hello messages are sent periodically on all interfaces where LDP is enabled.

There are two types of LDP hello messages:

- Link hello messages—Sent through the LDP interface as UDP packets addressed to the LDP discovery port. Receipt of an LDP link hello message on an interface identifies an adjacency with the LDP peer router.
- Targeted hello messages—Sent as UDP packets addressed to the LDP discovery port at a specific address. Targeted hello messages are used to support LDP sessions between routers that are not directly connected. A targeted router determines whether to respond or ignore a targeted hello message. A targeted router that chooses to

respond does so by periodically sending targeted hello messages back to the initiating router.

By default, LDP sends hello messages every 5 seconds for link hello messages and every 15 seconds for targeted hello messages. You can configure the LDP timer to alter how often both types of hello messages are sent. However, you cannot configure a time for the LDP timer that is greater than the LDP hold time. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 26](#).

### Configuring the LDP Timer for Link Hello Messages

To modify how often LDP sends link hello messages, specify a new link hello message interval for the LDP timer using the **hello-interval** statement:

```
hello-interval seconds;
```

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

### Configuring the LDP Timer for Targeted Hello Messages

To modify how often LDP sends targeted hello messages, specify a new targeted hello message interval for the LDP timer by configuring the **hello-interval** statement as an option for the **targeted-hello** statement:

```
targeted-hello {  
  hello-interval seconds;  
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

## Configuring the Delay Before LDP Neighbors Are Considered Down

---

The hold time determines how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. The values sent by each neighbor do not have to match.

The hold time should normally be at least three times the hello interval. The default is 15 seconds for link hello messages and 45 seconds for targeted hello messages. However, it is possible to configure an LDP hold time that is close to the value for the hello interval.



**NOTE:** By configuring an LDP hold time close to the hello interval (less than three times the hello interval), LDP neighbor failures might be detected more quickly. However, this also increases the possibility that the router might declare an LDP neighbor down that is still functioning normally. For more information, see [“Configuring the LDP Timer for Hello Messages” on page 25](#).

The LDP hold time is also negotiated automatically between LDP peers. When two LDP peers advertise different LDP hold times to one another, the smaller value is used. If an

LDP peer router advertises a shorter hold time than the value you have configured, the peer router's advertised hold time is used. This negotiation can affect the LDP keepalive interval as well.

If the local LDP hold time is not shortened during LDP peer negotiation, the user-configured keepalive interval is left unchanged. However, if the local hold time is reduced during peer negotiation, the keepalive interval is recalculated. If the LDP hold time has been reduced during peer negotiation, the keepalive interval is reduced to one-third of the new hold time value. For example, if the new hold-time value is 45 seconds, the keepalive interval is set to 15 seconds.

This automated keepalive interval calculation can cause different keepalive intervals to be configured on each peer router. This enables the routers to be flexible in how often they send keepalive messages, because the LDP peer negotiation ensures they are sent more frequently than the LDP hold time.

When you reconfigure the hold-time interval, changes do not take effect until after the session is reset. The hold time is negotiated when the LDP peering session is initiated and cannot be renegotiated as long as the session is up (required by RFC 5036, *LDP Specification*). To manually force the LDP session to reset, issue the **clear ldp session** command.

## Configuring the LDP Hold Time for Link Hello Messages

To modify how long an LDP node should wait for a link hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement:

```
hold-time seconds;
```

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

## Configuring the LDP Hold Time for Targeted Hello Messages

To modify how long an LDP node should wait for a targeted hello message before declaring the neighbor down, specify a new time in seconds using the **hold-time** statement as an option for the **targeted-hello** statement:

```
targeted-hello {  
  hold-time seconds;  
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary sections for these statements.

## Configuring the Interval for LDP Keepalive Messages

---

The keepalive interval determines how often a message is sent over the session to ensure that the keepalive timeout is not exceeded. If no other LDP traffic is sent over the session in this much time, a keepalive message is sent. The default is 10 seconds. The minimum value is 1 second.

The value configured for the keepalive interval can be altered during LDP session negotiation if the value configured for the LDP hold time on the peer router is lower than the value configured locally. For more information, see [“Configuring the Delay Before LDP Neighbors Are Considered Down” on page 26](#).

To modify the keepalive interval, include the **keepalive-interval** statement:

```
keepalive-interval seconds;
```

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

---

## Configuring the LDP Keepalive Timeout

After an LDP session is established, messages must be exchanged periodically to ensure that the session is still working. The keepalive timeout defines the amount of time that the neighbor LDP node waits before deciding that the session has failed. This value is usually set to at least three times the keepalive interval. The default is 30 seconds.

To modify the keepalive interval, include the **keepalive-timeout** statement:

```
keepalive-timeout seconds;
```

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

The value configured for the **keepalive-timeout** statement is displayed as the hold time when you issue the **show ldp session detail** command.

---

## Configuring LDP Route Preferences

When several protocols calculate routes to the same destination, route preferences are used to select which route is installed in the forwarding table. The route with the lowest preference value is selected. The preference value can be a number in the range 0 through 255. By default, LDP routes have a preference value of 9.

To modify the route preferences, include the **preference** statement:

```
preference preference;
```

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

---

## Configuring LDP Graceful Restart

When you alter the graceful restart configuration at either the **[edit routing-options graceful-restart]** or **[edit protocols ldp graceful-restart]** hierarchy levels, any running LDP session is automatically restarted to apply the graceful restart configuration. This behavior mirrors the behavior of BGP when you alter its graceful restart configuration.

By default, graceful restart helper mode is enabled, but graceful restart is disabled. Thus, the default behavior of a router is to assist neighboring routers attempting a graceful restart, but not to attempt a graceful restart itself.

To configure LDP graceful restart, see the following sections:

- [Enabling Graceful Restart on page 29](#)
- [Disabling LDP Graceful Restart or Helper Mode on page 29](#)
- [Configuring Reconnect Time on page 30](#)
- [Configuring Recovery Time and Maximum Recovery Time on page 30](#)

## Enabling Graceful Restart

To enable LDP graceful restart, you also need to enable graceful restart on the router.

To enable graceful restart, include the **graceful-restart** statement:

```
graceful-restart;
```

You can include this statement at the following hierarchy levels:

- **[edit routing-options]**
- **[edit logical-systems *logical-system-name* routing-options]**

The **graceful-restart** statement enables graceful restart for all protocols supporting this feature on the router. For more information about graceful restart, see the *Junos OS Routing Protocols Library for Routing Devices*.

By default, LDP graceful restart is enabled when you enable graceful restart at both the LDP protocol level and on all the routing instances. However, you can disable both LDP graceful restart and LDP graceful restart helper mode.

## Disabling LDP Graceful Restart or Helper Mode

To disable LDP graceful restart and recovery, include the **disable** statement:

```
ldp {
  graceful-restart {
    disable;
  }
}
```

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

You can disable helper mode at the LDP protocols level only. You cannot disable helper mode for a specific routing instance. To disable LDP helper mode, include the **helper-disable** statement:

```
ldp {
  graceful-restart {
    helper-disable;
  }
}
```

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

The following LDP graceful restart configurations are possible:

- LDP graceful restart and helper mode are both enabled.
- LDP graceful restart is disabled but helper mode is enabled. A router configured in this way cannot restart gracefully but can help a restarting neighbor.
- LDP graceful restart and helper mode are both disabled. The router does not use LDP graceful restart or the graceful restart type, length, and value (TLV) sent in the initialization message. The router behaves as a router that cannot support LDP graceful restart.

A configuration error is issued if you attempt to enable graceful restart and disable helper mode.

## Configuring Reconnect Time

After the LDP connection between neighbors fails, neighbors wait a certain amount of time for the gracefully restarting router to resume sending LDP messages. After the wait period, the LDP session can be reestablished. You can configure the wait period in seconds. This value is included in the fault tolerant session TLV sent in LDP initialization messages when LDP graceful restart is enabled.

Suppose that Router A and Router B are LDP neighbors. Router A is the restarting Router. The reconnect time is the time that Router A tells Router B to wait after Router B detects that Router A restarted.

To configure the reconnect time, include the **reconnect-time** statement:

```
graceful-restart {  
    reconnect-time seconds;  
}
```

You can set the reconnect time to a value in the range from 30 through 300 seconds. By default, it is 60 seconds.

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

## Configuring Recovery Time and Maximum Recovery Time

The recovery time is the amount of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This period is also typically the amount of time that a neighboring router maintains its information about the restarting router, allowing it to continue to forward traffic.

To prevent a neighboring router from being adversely affected if it receives a false value for the recovery time from the restarting router, you can configure the maximum recovery time on the neighboring router. A neighboring router maintains its state for the shorter of the two times. For example, Router A is performing an LDP graceful restart. It has sent a recovery time of 900 seconds to neighboring Router B. However, Router B has its maximum recovery time configured at 400 seconds. Router B will only wait for 400 seconds before it purges its LDP information from Router A.

To configure recovery time, include the **recovery-time** statement and the **maximum-neighbor-recovery-time** statement:

```

graceful-restart {
  maximum-neighbor-recovery-time seconds;
  recovery-time seconds;
}

```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

## Configuring the Prefixes Advertised into LDP from the Routing Table

You can control the set of prefixes that are advertised into LDP and cause the router to be the egress router for those prefixes. By default, only the loopback address is advertised into LDP. To configure the set of prefixes from the routing table to be advertised into LDP, include the **egress-policy** statement:

```
egress-policy policy-name;
```

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



**NOTE:** If you configure an egress policy for LDP that does not include the loopback address, it is no longer advertised in LDP. To continue to advertise the loopback address, you need to explicitly configure it as a part of the LDP egress policy.

The named policy (configured at the **[edit policy-options]** or **[edit logical-systems logical-system-name policy-options]** hierarchy level) is applied to all routes in the routing table. Those routes that match the policy are advertised into LDP. You can control the set of neighbors to which those prefixes are advertised by using the **export** statement. Only **from** operators are considered; you can use any valid **from** operator. For more information, see the *Junos OS Routing Protocols Library for Routing Devices*.

### Example: Configuring the Prefixes Advertised into LDP

Advertise all connected routes into LDP:

```

[edit protocols]
ldp {
  egress-policy connected-only;
}
policy-options {
  policy-statement connected-only {
    from {
      protocol direct;
    }
    then accept;
  }
}

```

## Configuring LDP LSP Traceroute

---

You can trace the route followed by an LDP-signaled LSP. LDP LSP traceroute is based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. This feature allows you to periodically trace all paths in a FEC. The FEC topology information is stored in a database accessible from the CLI.

A topology change does not automatically trigger a trace of an LDP LSP. However, you can manually initiate a traceroute. If the traceroute request is for an FEC that is currently in the database, the contents of the database are updated with the results.

The periodic traceroute feature applies to all FECs specified by the **oam** statement configured at the **[edit protocols ldp]** hierarchy level. To configure periodic LDP LSP traceroute, include the **periodic-traceroute** statement:

```
periodic-traceroute {  
  disable;  
  exp exp-value;  
  fanout fanout-value;  
  frequency minutes;  
  paths number-of-paths;  
  retries retry-attempts;  
  source address;  
  ttl ttl-value;  
  wait seconds;  
}
```

You can configure this statement at the following hierarchy levels:

- **[edit protocols ldp oam]**
- **[edit protocols ldp oam fec *address*]**

You can configure the **periodic-traceroute** statement by itself or with any of the following options:

- **exp**—Specify the class of service to use when sending probes.
- **fanout**—Specify the maximum number of next hops to search per node.
- **frequency**—Specify the interval between traceroute attempts.
- **paths**—Specify the maximum number of paths to search.
- **retries**—Specify the number of attempts to send a probe to a specific node before giving up.
- **source**—Specify the IPv4 source address to use when sending probes.
- **ttl**—Specify the maximum time-to-live value. Nodes that are beyond this value are not traced.
- **wait**—Specify the wait interval before resending a probe packet.



## Configuring Miscellaneous LDP Properties

---

The following sections describe how to configure a number of miscellaneous LDP properties:

- [Configuring LDP to Use the IGP Route Metric on page 33](#)
- [Preventing Addition of Ingress Routes to the inet.0 Routing Table on page 33](#)
- [Multiple-Instance LDP and Carrier-of-Carriers VPNs on page 34](#)
- [Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router on page 34](#)
- [Enabling LDP over RSVP-Established LSPs on page 34](#)
- [Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks on page 35](#)
- [Configuring the TCP MD5 Signature for LDP Sessions on page 35](#)
- [Configuring LDP Session Protection on page 36](#)
- [Disabling SNMP Traps for LDP on page 37](#)
- [Configuring LDP Synchronization with the IGP on LDP Links on page 37](#)
- [Configuring LDP Synchronization with the IGP on the Router on page 38](#)
- [Configuring the Label Withdrawal Timer on page 38](#)
- [Ignoring the LDP Subnet Check on page 38](#)

### Configuring LDP to Use the IGP Route Metric

Use the **track-igp-metric** statement if you want the interior gateway protocol (IGP) route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).

To use the IGP route metric, include the **track-igp-metric** statement:

```
track-igp-metric;
```

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

### Preventing Addition of Ingress Routes to the inet.0 Routing Table

By configuring the **no-forwarding** statement, you can prevent ingress routes from being added to the inet.0 routing table instead of the inet.3 routing table even if you enabled the **traffic-engineering bgp-igp** statement at the **[edit protocols mpls]** or the **[edit logical-systems *logical-system-name* protocols mpls]** hierarchy level. By default, the **no-forwarding** statement is disabled.

To omit ingress routes from the inet.0 routing table, include the **no-forwarding** statement:

```
no-forwarding;
```

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

## Multiple-Instance LDP and Carrier-of-Carriers VPNs

By configuring multiple LDP routing instances, you can use LDP to advertise labels in a carrier-of-carriers VPN from a service provider provider edge (PE) router to a customer carrier customer edge (CE) router. This is especially useful when the carrier customer is a basic Internet service provider (ISP) and wants to restrict full Internet routes to its PE routers. By using LDP instead of BGP, the carrier customer shields its other internal routers from the Internet. Multiple-instance LDP is also useful when a carrier customer wants to provide Layer 2 or Layer 3 VPN services to its customers.

For an example of how to configure multiple LDP routing instances for carrier-of-carriers VPNs, see the *Multiple Instances for Label Distribution Protocol Feature Guide*.

## Configuring MPLS and LDP to Pop the Label on the Ultimate-Hop Router

The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop router removes the label and sends the packet to the egress router. If ultimate-hop popping is enabled, label 0 (IPv4 Explicit Null label) is advertised. Ultimate-hop popping ensures that any packets traversing an MPLS network include a label.

To configure ultimate-hop popping, include the **explicit-null** statement:

```
explicit-null;
```

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



**NOTE:** Juniper Networks routers queue packets based on the incoming label. Routers from other vendors might queue packets differently. Keep this in mind when working with networks containing routers from multiple vendors.

For more information about labels, see *MPLS Label Overview* and *MPLS Label Allocation*.

## Enabling LDP over RSVP-Established LSPs

You can run LDP over LSPs established by RSVP, effectively tunneling the LDP-established LSP through the one established by RSVP. To do so, enable LDP on the lo0.0 interface (see “[Enabling and Disabling LDP](#)” on page 10). You must also configure the LSPs over which you want LDP to operate by including the **ldp-tunneling** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level:

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-name {
      from source;
      to destination;
      ldp-tunneling;
    }
  }
}
```

```
}
```

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

**Related  
Documentation**

- [Tunneling LDP LSPs in RSVP LSPs Overview on page 5](#)

## Enabling LDP over RSVP-Established LSPs in Heterogeneous Networks

Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This might require that you manually configure the RSVP metric when deploying LDP tunneling over RSVP LSPs in heterogeneous networks.

When a Juniper Networks router is linked to another vendor's router through an RSVP tunnel, and LDP tunneling is also enabled, by default the Juniper Networks router might not use the RSVP tunnel to route traffic to the LDP destinations downstream of the other vendor's egress router if the RSVP path has a metric of 1 larger than the physical OSPF path.

To ensure that LDP tunneling functions properly in heterogeneous networks, you can configure OSPF to ignore the RSVP LSP metric by including the **ignore-lsp-metrics** statement:

```
ignore-lsp-metrics;
```

You can configure this statement at the following hierarchy levels:

- [\[edit protocols ospf traffic-engineering shortcuts\]](#)
- [\[edit logical-systems \*logical-system-name\* protocols ospf traffic-engineering shortcuts\]](#)

To enable LDP over RSVP LSPs, you also still need to complete the procedure in Section [“Enabling LDP over RSVP-Established LSPs” on page 34](#).

## Configuring the TCP MD5 Signature for LDP Sessions

You can configure an MD5 signature for an LDP TCP connection to protect against the introduction of spoofed TCP segments into LDP session connection streams.

A router using the MD5 signature option is configured with a password for each peer for which authentication is required. The password is stored encrypted.

LDP hello adjacencies can still be created even when peering interfaces are configured with different security signatures. However, the TCP session cannot be authenticated and is never established.

To configure an MD5 signature for an LDP TCP connection, include the **session** and **authentication-key** statement:

```
session address {
  authentication-key md5-authentication-key;
}
```

For a list of hierarchy levels at which you can include these statements, see the statement summary section for the **session** statement.

Use the **session** statement to configure the address for the remote end of the LDP session.

The **md5-authentication-key** (password) can be up to 69 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks.

You can also configure an authentication key update mechanism for the LDP routing protocol. This mechanism allows you to update authentication keys without interrupting associated routing and signaling protocols such as Open Shortest Path First (OSPF) and Resource Reservation Setup Protocol (RSVP).

To configure the authentication key update mechanism, include the **key-chain** statement at the **[edit security authentication-key-chains]** hierarchy level, and specify the **key** option to create a keychain consisting of several authentication keys.

```
[edit security authentication-key-chains]
key-chain key-chain-name {
  key key {
    secret secret-data;
    start-time yyyy-mm-dd.hh:mm:ss;
  }
}
```

To configure the authentication key update mechanism for the LDP routing protocol, include the **authentication-key-chain** statement at the **[edit protocols ldp]** hierarchy level to associate the protocol with the **[edit security authentication-key-chains]** authentication keys.

```
[edit protocols ldp]
group group-name {
  neighbor address {
    authentication-key-chain key-chain-name;
  }
}
```

For more information about the authentication key update feature, see *Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols*.

## Configuring LDP Session Protection

An LDP session is normally created between a pair of routers that are connected by one or more links. The routers form one hello adjacency for every link that connects them and associate all the adjacencies with the corresponding LDP session. When the last hello adjacency for an LDP session goes away, the LDP session is terminated. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished.

You can configure the Junos OS to leave the LDP session between two routers up even if there are no hello adjacencies on the links connecting the two routers by configuring the **session-protection** statement. You can optionally specify a time in seconds using the **timeout** option. The session remains up for the duration specified as long as the routers maintain IP network connectivity.

```
session-protection {
```

```

    timeout seconds;
}

```

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

## Disabling SNMP Traps for LDP

Whenever an LDP LSP makes a transition from up to down, or down to up, the router sends an SNMP trap. However, it is possible to disable the LDP SNMP traps on a router, logical system, or routing instance.

For information about the LDP SNMP traps and the proprietary LDP MIB, see the *SNMP MIBs and Traps Reference* and *Interpreting the Enterprise-Specific LDP MIB*.

To disable SNMP traps for LDP, specify the **trap disable** option for the **log-updown** statement:

```

log-updown {
    trap disable;
}

```

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

## Configuring LDP Synchronization with the IGP on LDP Links

LDP is a protocol for distributing labels in non-traffic-engineered applications. Labels are distributed along the best path determined by the IGP. If synchronization between LDP and the IGP is not maintained, the LSP goes down. When LDP is not fully operational on a given link (a session is not established and labels are not exchanged), the IGP advertises the link with the maximum cost metric. The link is not preferred but remains in the network topology.

LDP synchronization is supported only on active point-to-point interfaces and LAN interfaces configured as point-to-point under the IGP. LDP synchronization is not supported during graceful restart.

To advertise the maximum cost metric until LDP is operational for synchronization, include the **ldp-synchronization** statement:

```

ldp-synchronization {
    disable;
    hold-time seconds;
}

```

To disable synchronization, include the **disable** statement. To configure the time period to advertise the maximum cost metric for a link that is not fully operational, include the **hold-time** statement.

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

## Configuring LDP Synchronization with the IGP on the Router

You can configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.

To configure the time the LDP waits before informing the IGP that the LDP neighbor and session are operational, include the **igp-synchronization** statement and specify a time in seconds for the **holddown-interval** option:

```
igp-synchronization holddown-interval seconds;
```

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

## Configuring the Label Withdrawal Timer

The label withdrawal timer delays sending a label withdrawal message for a FEC to a neighbor. When an IGP link to a neighbor fails, the label associated with the FEC has to be withdrawn from all the upstream routers if the neighbor is the next hop for the FEC. After the IGP converges and a label is received from a new next hop, the label is readvertised to all the upstream routers. This is the typical network behavior. By delaying label withdrawal by a small amount of time (for example, until the IGP converges and the router receives a new label for the FEC from the downstream next hop), the label withdrawal and sending a label mapping soon could be avoided. The **label-withdrawal-delay** statement allows you to configure this delay time. By default, the delay is 60 seconds.

If the router receives the new label before the timer runs out, the label withdrawal timer is canceled. However, if the timer runs out, the label for the FEC is withdrawn from all of the upstream routers.

By default, LDP waits for 60 seconds before withdrawing labels to avoid resignaling LSPs multiple times while the IGP is reconverging. To configure the label withdrawal delay time in seconds, include the **label-withdrawal-delay** statement:

```
label-withdrawal-delay seconds;
```

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

## Ignoring the LDP Subnet Check

In Junos OS Release 8.4 and later releases, an LDP source address subnet check is performed during the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address. This causes an interoperability issue with some other vendors' equipment.

To disable the subnet check, include the **allow-subnet-mismatch** statement:

```
allow-subnet-mismatch;
```

This statement can be included at the following hierarchy levels:

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





## CHAPTER 2

# Configuration Statements for LDP

- [allow-subnet-mismatch](#) on page 42
- [authentication-algorithm](#) on page 43
- [authentication-key \(Protocols LDP\)](#) on page 45
- [authentication-key-chain \(Protocols LDP\)](#) on page 46
- [deaggregate](#) on page 47
- [disable \(Protocols LDP\)](#) on page 48
- [dod-request-policy](#) on page 49
- [downstream-on-demand](#) on page 49
- [ecmp](#) on page 50
- [egress-policy](#) on page 50
- [explicit-null \(Protocols LDP\)](#) on page 51
- [export \(Protocols LDP\)](#) on page 51
- [fec](#) on page 52
- [graceful-restart \(Protocols LDP\)](#) on page 53
- [hello-interval \(Protocols LDP\)](#) on page 54
- [helper-disable \(LDP\)](#) on page 55
- [hold-time \(Protocols LDP\)](#) on page 56
- [ignore-lsp-metrics](#) on page 57
- [igp-synchronization](#) on page 57
- [import \(Protocols LDP\)](#) on page 58
- [interface \(Protocols LDP\)](#) on page 59
- [keepalive-interval](#) on page 60
- [keepalive-timeout](#) on page 61
- [l2-smart-policy](#) on page 61
- [label-withdrawal-delay](#) on page 62
- [ldp](#) on page 63
- [ldp-synchronization](#) on page 66
- [ldp-tunneling](#) on page 66

- [log-updown \(Protocols LDP\) on page 67](#)
- [maximum-neighbor-recovery-time on page 68](#)
- [no-forwarding on page 69](#)
- [policing \(Protocols LDP\) on page 70](#)
- [preference \(Protocols LDP\) on page 71](#)
- [reconnect-time on page 72](#)
- [recovery-time on page 73](#)
- [session \(ldp\) on page 74](#)
- [session-protection on page 75](#)
- [strict-targeted-hellos on page 75](#)
- [targeted-hello on page 76](#)
- [traceoptions \(Protocols LDP\) on page 77](#)
- [track-igp-metric on page 79](#)
- [traffic-statistics \(Protocols LDP\) on page 80](#)
- [transport-address on page 82](#)

---

## allow-subnet-mismatch

---

<b>Syntax</b>	<code>allow-subnet-mismatch;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i> ], [edit protocols ldp interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.3. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Ignore the LDP subnet check. For Junos OS Release 8.4 and later releases, an LDP source address subnet check was added for the neighbor establishment procedure. The source address in the LDP link hello packet is matched against the interface address.
<b>Default</b>	The source address in the LDP link hello packet is matched against the interface 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="#">Ignoring the LDP Subnet Check on page 38</a></li></ul>

## authentication-algorithm

<b>Syntax</b>	<code>authentication-algorithm <i>algorithm</i>;</code>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> protocols ldp session <i>session-address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols   ldp session <i>session-address</i>], [edit logical-systems <i>logical-system-name</i> routing-options bmp], [edit logical-systems <i>logical-system-name</i> routing-options bmp station <i>station-name</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols ldp session <i>session-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>   neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols ldp session <i>session-address</i>], [edit routing-options bmp], [edit routing-options bmp station <i>station-name</i>]</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 7.6.</p> <p>Statement introduced for BGP in Junos OS Release 8.0.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p> <p>Statement introduced for BMP in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced for BMP in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Configure an authentication algorithm type.



**NOTE:** Keep the following points in mind when you configure the authentication algorithm in an IPsec proposal:

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, an error occurs and the tunnel is not established in this scenario. For example, if one end of the tunnel contains router 1 configured with the authentication algorithm as hmac-sha-256-128 and the other end of the tunnel contains router 2 configured with the authentication algorithm as hmac-md5-96, the VPN tunnel is not established.

- When both ends of an IPsec VPN tunnel contain the same IKE proposal but different IPsec proposals, and when one end of the tunnel contains two IPsec proposals to check whether a less secure algorithm is selected or not, an error occurs and the tunnel is not established. For example, if you configure two authentication algorithms for an IPsec proposal as hmac-sha-256-128 and hmac-md5-96 on one end of the tunnel, router 1, and if you configure the algorithm for an IPsec proposal as hmac-md5-96 on the other end of the tunnel, router 2, the tunnel is not established and the number of proposals mismatch.
  - When you configure two IPsec proposals at both ends of a tunnel, such as the authentication-algorithm hmac-sha-256-128 and authentication-algorithm hmac-md5-96 statements at the [edit services ipsec-vpn ipsec proposal *proposal-name*] hierarchy level on one of the tunnel, router 1 (with the algorithms in two successive statements to specify the order), and the authentication-algorithm hmac-md5-96 and authentication-algorithm hmac-sha-256-128 statements at the [edit services ipsec-vpn ipsec proposal *proposal-name*] hierarchy level on one of the tunnel, router 2 (with the algorithms in two successive statements to specify the order, which is the reverse order of router 1), the tunnel is established in this combination as expected because the number of proposals is the same on both ends and they contain the same set of algorithms. However, the authentication algorithm selected is hmac-md5-96 and not the stronger algorithm of hmac-sha-256-128. This method of selection of the algorithm occurs because the first matching proposal is selected. Also, for a default proposal, regardless of whether the router supports the Advanced Encryption Standard (AES) encryption algorithm, the 3des-cbc algorithm is chosen and not the aes-cfb algorithm, which is because of the first algorithm in the default proposal being selected. In the sample scenario described here, on router 2, if you reverse the order of the algorithm configuration in the proposal so that it is the same order as the one specified on router 1, hmac-sha-256-128 is selected as the authentication method.
  - You must be aware of the order of proposals in an IPsec policy at the time of configuration if you want the matching of proposals to happen in a certain order of preference, such as the strongest algorithm to be considered first when a match is made when both policies from the two peers have a proposal.
-

**Options** *algorithm*—Specify one of the following types of authentication algorithms:

- **aes-128-cmac-96**—Cipher-based message authentication code (AES128, 96 bits).
- **hmac-sha-1-96**—Hash-based message authentication code (SHA1, 96 bits).
- **md5**—Message digest 5.

**Default:** hmac-sha-1-96



**NOTE:** The default is not displayed in the output of the `show bgp bmp` command unless a key or key-chain is also configured.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Example: Configuring Route Authentication for BGP*
- *Configuring BGP Monitoring Protocol Version 3*

## authentication-key (Protocols LDP)

**Syntax** authentication-key *md5-authentication-key*;

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols ldp session *address*],  
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ldp session *address*],  
[edit protocols ldp session *address*],  
[edit routing-instances *routing-instance-name* protocols ldp session *address*]

**Release Information** Statement introduced before Junos OS Release 7.4.  
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Configure the MD5 authentication signature. The maximum length of the authentication signature is 69 characters.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring the TCP MD5 Signature for LDP Sessions on page 35](#)

## authentication-key-chain (Protocols LDP)

---

<b>Syntax</b>	authentication-key-chain <i>key-chain</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>name</i> protocols ldp session <i>address</i> ], [edit logical-systems <i>name</i> routing-instances <i>instance-name</i> protocols ldp session <i>address</i> ], [edit protocols ldp session <i>address</i> ], [edit routing-instances <i>instance-name</i> protocols ldp session <i>address</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Apply and enable an authentication keychain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for LDP, you cannot commit the <b>0.0.0.0/allow</b> statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations.
<b>Options</b>	<b>key-chain</b> —Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
<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 Authentication Key Update Mechanism for BGP and LDP Routing Protocols</a></li><li>• <a href="#">Configuring Miscellaneous LDP Properties on page 33</a></li></ul>

## deaggregate

---

<b>Syntax</b>	deaggregate   no-deaggregate;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Control forwarding equivalence class (FEC) deaggregation on the router. The use of the <b>deaggregate</b> statement in LDP is a standard practice that we recommend for LDP deployments.
<b>Default</b>	Deaggregation is disabled on the router.
<b>Options</b>	<b>deaggregate</b> —Deaggregate FECs. <b>no-deaggregate</b> —Aggregate FECs.
<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>• <i>Configuring FEC Deaggregation</i></li> </ul>

## disable (Protocols LDP)

---

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options graceful-restart], [edit protocols ldp graceful-restart], [edit protocols ldp interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> routing-options graceful-restart]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Explicitly disable LDP on an interface, or explicitly disable LDP graceful restart.
<b>Default</b>	LDP is enabled on interfaces configured with the LDP <b>interface</b> statement. LDP graceful restart is automatically enabled when graceful restart is enabled under the <b>[edit routing-options]</b> 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="#">Enabling and Disabling LDP on page 10</a></li><li>• <a href="#">Configuring LDP Graceful Restart on page 28</a></li></ul>



## dod-request-policy

---

<b>Syntax</b>	<code>dod-request-policy <i>dod-request-policy-name</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit protocols ldp]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the name of the LDP downstream on demand request policy. LDP sends label request messages only for those FECs matching in the downstream on demand request policy.
<b>Options</b>	<i>dod-request-policy-name</i> —Specify the name of the downstream on demand request 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="#">Example: Configuring LDP Downstream on Demand on page 20</a></li> </ul>

## downstream-on-demand

---

<b>Syntax</b>	<code>downstream-on-demand;</code>
<b>Hierarchy Level</b>	[edit logical systems <i>logical-system-name</i> protocols ldp session <i>session-address</i> ], [edit protocols ldp session <i>session-address</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 12.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable LDP downstream on demand on the LDP session. LDP is widely deployed in downstream unsolicited advertisement mode. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between access and aggregation networks to minimize the workload for the access node (AN) control plane and to avoid the storage of tens of thousands of label bindings from upstream aggregation nodes.
<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 LDP Downstream on Demand on page 20</a></li> </ul>

## ecmp

<b>Syntax</b>	<code>ecmp;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp oam bfd-liveness-detection], [edit logical-systems <i>logical-system-name</i> protocols ldp oam fec <i>address</i> bfd-liveness-detection], [edit protocols ldp oam bfd-liveness-detection], [edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.
<b>Description</b>	Allows LDP to establish BFD sessions for all ECMP paths configured for the specified FEC. If you configure the <b>ecmp</b> statement, you must also configure the <b>periodic-traceroute</b> statement for the specified FEC. If you do not do so, the commit operation fails. You can configure the <b>periodic-traceroute</b> statement at the global hierarchy level ([edit protocols ldp oam]) while only configuring the <b>ecmp</b> statement for a specific FEC ([edit protocols ldp oam fec <i>address</i> bfd-liveness-detection]).
<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 ECMP-Aware BFD for LDP LSPs</a></li> </ul>

## egress-policy

<b>Syntax</b>	<code>egress-policy [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Control the prefixes advertised into LDP.
<b>Default</b>	Only the loopback address is advertised.
<b>Options</b>	<i>policy-names</i> —Name of one or more routing policies.
<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 Prefixes Advertised into LDP from the Routing Table on page 31</a></li> </ul>

## explicit-null (Protocols LDP)

<b>Syntax</b>	<code>explicit-null;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Advertise label 0 to the egress router of a label-switched path (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 MPLS and LDP to Pop the Label on the Ultimate-Hop Router on page 34</a></li> </ul>

## export (Protocols LDP)

<b>Syntax</b>	<code>export [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Apply policy filters to outbound LDP label bindings. Filters are applied to all label bindings from all neighbors.
<b>Options</b>	<i>policy-names</i> —Name of one or more routing policies.
<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="#">Filtering Outbound LDP Label Bindings on page 13</a></li> </ul>

## fec

<b>Syntax</b>	<pre> fec <i>fec-address</i> {     bfd-liveness-detection {         detection-time threshold <i>milliseconds</i>;         ecmp;         failure-action {             remove-nexthop;             remove-route;         }         holddown-interval <i>milliseconds</i>;         ingress-policy <i>ingress-policy-name</i>;         minimum-interval <i>milliseconds</i>;         minimum-receive-interval <i>milliseconds</i>;         minimum-transmit-interval <i>milliseconds</i>;         multiplier <i>detection-time-multiplier</i>;         no-adaptation;         transmit-interval {             minimum-interval <i>milliseconds</i>;             threshold <i>milliseconds</i>;         }         version (0   1   automatic);     }     no-bfd-liveness-detection;     periodic-traceroute {         disable;         exp <i>exp-value</i>;         fanout <i>fanout-value</i>;         frequency <i>minutes</i>;         paths <i>number-of-paths</i>;         retries <i>retry-attempts</i>;         source <i>address</i>;         ttl <i>ttl-value</i>;         wait <i>seconds</i>;     } } </pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-systems-name</i> protocols ldp oam], [edit protocols ldp oam]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 12.2 for EX Series switches. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Allows you to configure BFD for a specific LDP forwarding equivalence class (FEC).
<b>Options</b>	<p><b><i>fec-address</i></b>—Specify the FEC address.</p> <p>The other 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>

**Related Documentation** • [Configuring BFD for LDP LSPs](#)

## graceful-restart (Protocols LDP)

<b>Syntax</b>	<pre>graceful-restart {   disable;   helper-disable;   maximum-neighbor-recovery-time <i>value</i>;   reconnect-time <i>seconds</i>;   recovery-time <i>value</i>; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	Configure LDP graceful restart on the LDP master protocol instance or for a specific routing instance.



**NOTE:** When you alter the graceful restart configuration at either the [edit routing-options graceful-restart] or [edit protocols ldp graceful-restart] hierarchy levels, any running LDP session is automatically restarted to apply the graceful restart configuration. This behavior mirrors the behavior of BGP when you alter its graceful restart configuration.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation** • [Configuring LDP Graceful Restart on page 28](#)

## hello-interval (Protocols LDP)

---

<b>Syntax</b>	<code>hello-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>  ldp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>  ldp targeted-hello],</code> <code>[edit protocols ldp interface <i>interface-name</i>],</code> <code>[edit protocols ldp targeted-hello],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Support for LDP targeted hellos added in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Control the LDP timer that regulates how often hello messages are sent. You can control the rate both link hello messages and targeted hello messages are sent depending on the hierarchy level at which you configure the <b>hello-interval</b> statement.
<b>Options</b>	<b><i>seconds</i></b> —Length of time between transmission of hello packets. <b>Range:</b> 1 through 65,535 seconds <b>Default:</b> 5 seconds for link hello messages, 15 seconds for targeted hello messages
<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 LDP Timer for Hello Messages on page 25</a></li></ul>

---

## helper-disable (LDP)

---

<b>Syntax</b>	helper-disable;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Disable helper mode for LDP graceful restart. When helper mode is disabled, a router cannot help a neighboring router that is attempting to restart LDP.
<b>Default</b>	Helper mode is enabled by default on all routing protocols (including LDP) that support graceful restart.
<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 LDP Graceful Restart on page 28</a></li></ul>

## hold-time (Protocols LDP)

---

<b>Syntax</b>	<code>hold-time seconds;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols ldp targeted-hello],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>  ldp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols</code> <code>  ldp targeted-hello],</code> <code>[edit protocols ldp interface <i>interface-name</i>],</code> <code>[edit protocols ldp targeted-hello],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> protocols ldp targeted-hello]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Support for LDP targeted hellos added in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify how long an LDP node should wait for a hello message before declaring a neighbor to be down. This value is sent as part of a hello message so that each LDP node tells its neighbors how long to wait. You can specify times for both link hello messages and targeted hello messages depending on the hierarchy level at which you configure the <b>hold-time</b> statement.
<b>Options</b>	<b>seconds</b> —Hold-time value. <b>Range:</b> 1 through 65,535 seconds <b>Default:</b> 15 seconds for link hello messages, 45 seconds for targeted hello messages
<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 Delay Before LDP Neighbors Are Considered Down on page 26</a></li></ul>



## ignore-lsp-metrics

<b>Syntax</b>	ignore-lsp-metrics;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ospf traffic-engineering shortcuts], [edit protocols ospf traffic-engineering shortcuts]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Cause OSPF to ignore the RSVP LSP metric.  Some other vendors use an OSPF metric of 1 for the loopback address. Juniper Networks routers use an OSPF metric of 0 for the loopback address. This can cause interoperability problems when you configure LDP tunneling over RSVP LSPs in heterogeneous networks.
<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 LDP over RSVP-Established LSPs in Heterogeneous Networks on page 35</a></li> </ul>

## igp-synchronization

<b>Syntax</b>	igp-synchronization holddown-interval <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Configure the time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. For large networks with numerous FECs, you might need to configure a longer value to allow enough time for the LDP label databases to be exchanged.
<b>Options</b>	<b>holddown-interval <i>seconds</i></b> —Time the LDP waits before informing the IGP that the LDP neighbor and session for an interface are operational. <b>Default:</b> 10 seconds <b>Range:</b> 10 through 60 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 LDP Synchronization with the IGP on the Router on page 38</a></li> </ul>

## import (Protocols LDP)

---

<b>Syntax</b>	<code>import [ <i>policy-names</i> ];</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Apply policy filters to received LDP label bindings. Filters are applied to all label bindings from all neighbors.
<b>Options</b>	<i>policy-names</i> —Name of one or more routing policies.
<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="#">Filtering Inbound LDP Label Bindings on page 11</a></li></ul>

## interface (Protocols LDP)

<b>Syntax</b>	<pre>interface <i>interface-name</i> {     disable;     hello-interval <i>seconds</i>;     hold-time <i>seconds</i>;     transport-address (interface   loopback); }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	Enable LDP on one or more router interfaces.
<b>Default</b>	LDP is disabled on all interfaces.
<b>Options</b>	<p><i>interface-name</i>—Name of an interface. To configure all interfaces, specify <b>all</b>.</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="#">Enabling and Disabling LDP on page 10</a></li> </ul>

## keepalive-interval

---

<b>Syntax</b>	<code>keepalive-interval <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Set the keepalive interval value.
<b>Options</b>	<b><i>seconds</i></b> —Keepalive value. <b>Range:</b> 1 through 65,535 <b>Default:</b> 10 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 Interval for LDP Keepalive Messages on page 27</a></li></ul>

## keepalive-timeout

<b>Syntax</b>	<code>keepalive-timeout <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Set the keepalive timeout value. The keepalive timeout defines the amount of time that the neighbor LDP node waits before determining that the session has failed.
<b>Options</b>	<b><i>seconds</i></b> —Keepalive timeout value. <b>Range:</b> 1 through 65,535 <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 LDP Keepalive Timeout on page 28</a></li> </ul>

## l2-smart-policy

<b>Syntax</b>	<code>l2-smart-policy;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Prevent LDP from exporting IPv4 FECs over sessions with Layer 2 neighbors only. IPv4 FECs received over such sessions are filtered out.
<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 LDP IPv4 FEC Filtering</a></li> </ul>

## label-withdrawal-delay

---

<b>Syntax</b>	label-withdrawal-delay <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Delay the withdrawal of labels to reduce router workload during IGP convergence.
<b>Options</b>	<b>seconds</b> —Configure the number of seconds to wait before withdrawing labels for the LDP LSPs. <b>Default:</b> 60 seconds <b>Range:</b> 0 through 300 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 Label Withdrawal Timer on page 38</a></li></ul>

## ldp

```
Syntax  ldp {
    (deaggregate | no-deaggregate);
    egress-policy [ policy-names ];
    explicit-null;
    export [ policy-names ];
    graceful-restart {
        disable;
        helper-disable;
        maximum-neighbor-recovery-time seconds;
        reconnect-time seconds;
        recovery-time seconds;
    }
    import [ policy-names ];
    interface (interface-name | all) {
        disable;
        hello-interval seconds;
        hold-time seconds;
        transport-address (interface | router-id);
    }
    keepalive-interval seconds;
    keepalive-timeout seconds;
    log-updown {
        trap disable;
    }
    no-forwarding;
    oam {
        bfd-liveness-detection {
            detection-time threshold milliseconds;
            ecmp;
            failure-action {
                remove-nexthop;
                remove-route;
            }
            holddown-interval milliseconds;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-transmit-interval milliseconds;
            multiplier detection-time-multiplier;
            no-adaptation;
            transmit-interval {
                minimum-interval milliseconds;
                threshold milliseconds;
            }
        }
    }
    fec fec-address {
        bfd-liveness-detection {
            detection-time threshold milliseconds;
            ecmp;
            failure-action {
                remove-nexthop;
                remove-route;
            }
        }
    }
}
```

```

        holddown-interval milliseconds;
        ingress-policy ingress-policy-name;
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        minimum-transmit-interval milliseconds;
        multiplier detection-time-multiplier;
        no-adaptation;
        transmit-interval {
            minimum-interval milliseconds;
            threshold milliseconds;
        }
        version (0 | 1 | automatic);
    }
    no-bfd-liveness-detection;
    periodic-traceroute {
        disable;
        exp exp-value;
        fanout fanout-value;
        frequency minutes;
        paths number-of-paths;
        retries retry-attempts;
        source address;
        ttl ttl-value;
        wait seconds;
    }
}
ingress-policy ingress-policy-name;
periodic-traceroute {
    disable;
    exp exp-value;
    fanout fanout-value;
    frequency minutes;
    paths number-of-paths;
    retries retry-attempts;
    source address;
    ttl ttl-value;
    wait seconds;
}
}
p2mp;
policing {
    fec fec-address {
        ingress-traffic filter-name;
        transit-traffic filter-name;
    }
}
preference preference;
session address {
    authentication-algorithm algorithm;
    authentication-key authentication-key;
    authentication-key-chain key-chain-name;
}
strict-targeted-hellos;
traceoptions {
    file filename <files number <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;

```



```

}
track-igp-metric;
traffic-statistics {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  interval interval;
  no-penultimate-hop;
}
transport-address (address | interface | router-id);
}

```

<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for EX Series switches. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable LDP routing on the router or switch.  You must include the <b>ldp</b> statement in the configuration to enable LDP on the router or switch.
<b>Default</b>	LDP is disabled on the router.
<b>Options</b>	The other 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="#">Minimum LDP Configuration on page 10</a></li> <li>• <a href="#">Enabling and Disabling LDP on page 10</a></li> </ul>

## ldp-synchronization

---

<b>Syntax</b>	<code>ldp-synchronization {     disable;     hold-time seconds; }</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ospf interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ospf interface <i>interface-name</i> ], [edit protocols ospf interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols ospf interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable synchronization by advertising the maximum cost metric until LDP is operational on the link.
<b>Options</b>	The other 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 LDP Synchronization with the IGP on LDP Links on page 37</a></li></ul>

## ldp-tunneling

---

<b>Syntax</b>	<code>ldp-tunneling;</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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable the LSP to be used for LDP tunneling.
<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 LDP over RSVP-Established LSPs on page 34</a></li></ul>

---

## log-updown (Protocols LDP)

---

<b>Syntax</b>	log-updown { trap disable; }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Disable LDP traps on the router, logical system, or routing instance.
<b>Options</b>	<b>trap disable</b> —Disable LDP traps. <b>Default:</b> LDP traps are enabled 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="#">Disabling SNMP Traps for LDP on page 37</a></li></ul>

## maximum-neighbor-recovery-time

---

<b>Syntax</b>	<code>maximum-neighbor-recovery-time seconds;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement changed from <b>maximum-recovery-time</b> to <b>maximum-neighbor-recovery-time</b> in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the maximum amount of time to wait before giving up an attempt to gracefully restart.
<b>Options</b>	<b>seconds</b> —Configure the maximum recovery time, in seconds. <b>Range:</b> 120 through 1800 seconds <b>Default:</b> 140 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 Recovery Time and Maximum Recovery Time on page 30</a></li><li>• <a href="#">Configuring Graceful Restart Options for LDP</a></li><li>• <a href="#">no-strict-lsa-checking</a></li><li>• <a href="#">recovery-time</a></li></ul>

## no-forwarding

---

<b>Syntax</b>	no-forwarding;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Do not add ingress routes to the inet.0 routing table even if <a href="#">traffic-engineering bgp-igp</a> (configured at the <a href="#">[edit protocols mpls]</a> hierarchy level) is enabled.
<b>Default</b>	The <b>no-forwarding</b> statement is disabled. Ingress routes are added to the inet.0 routing table instead of the inet.3 routing table when <a href="#">traffic-engineering bgp-igp</a> is 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="#">Preventing Addition of Ingress Routes to the inet.0 Routing Table on page 33</a></li> <li>• <a href="#">Configuring Virtual-Router Routing Instances in VPNs</a></li> </ul>

## policing (Protocols LDP)

---

<b>Syntax</b>	<pre>policing {     fec <i>fec-address</i> {         ingress-traffic <i>filter-name</i>;         transit-traffic <i>filter-name</i>;     } }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable policing of forwarding equivalence classes (FECs) for LDP.
<b>Options</b>	<p><b>fec <i>fec-address</i></b>—Specify the address for the FEC.</p> <p><b>ingress-traffic <i>filter-name</i></b>—Specify the name of the filter for policing ingress FEC traffic.</p> <p><b>transit-traffic <i>filter-name</i></b>—Specify the name of the filter for policing transit FEC traffic.</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>• <i>Configuring Policers for LDP FECs</i></li></ul>

## preference (Protocols LDP)

---

<b>Syntax</b>	<code>preference <i>preference</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit protocols ldp interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit protocols ldp interface <i>interface-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Set the route preference level for LDP routes.
<b>Options</b>	<i>preference</i> —Preferred value. <b>Range:</b> 0 through 255 <b>Default:</b> 9
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP Route Preferences on page 28</a></li> </ul>

## reconnect-time

---

<b>Syntax</b>	<code>reconnect-time seconds;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp <a href="#">graceful-restart</a> ], [edit protocols ldp <a href="#">graceful-restart</a> ], [edit routing-instances <i>routing-instance-name</i> protocols ldp <a href="#">graceful-restart</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the length of time required to reestablish a Label Distribution Protocol (LDP) session after graceful restart.
<b>Options</b>	<b>seconds</b> —Time required for reconnection. <b>Range:</b> 30 through 300 <b>Default:</b> 60 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 LDP Graceful Restart on page 28</a> on <i>MPLS Applications Feature Guide for Routing Devices</i></li><li>• <i>Configuring Graceful Restart Options for LDP</i></li></ul>



## recovery-time

---

<b>Syntax</b>	<code>recovery-time seconds;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp graceful-restart], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart], [edit protocols ldp graceful-restart], [edit routing-instances <i>routing-instance-name</i> protocols ldp graceful-restart]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the amount of time a router waits for LDP to restart gracefully.
<b>Options</b>	<b>seconds</b> —Configure the recovery time, in seconds. <b>Range:</b> 120 through 1800 seconds <b>Default:</b> 140 seconds
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Recovery Time and Maximum Recovery Time on page 30</a></li> </ul>

## session (ldp)

---

<b>Syntax</b>	<pre>session address {     authentication-algorithm <i>algorithm</i>;     authentication-key <i>authentication-key</i>;     authentication-key-chain <i>key-chain-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. <b>authentication-algorithm</b> statement introduced in Junos OS Release 7.6. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the address for the remote end of the LDP session.  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>• <a href="#">Configuring the TCP MD5 Signature for LDP Sessions on page 35</a></li></ul>

## session-protection

---

<b>Syntax</b>	session-protection { timeout <i>seconds</i> ; }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Description</b>	Configure when an LDP session is torn down and resigaled after the router stops receiving hello messages from a neighboring router. You might want to modify this behavior to prevent an LDP session from being unnecessarily terminated and reestablished. The LDP session remains up for the duration specified as long as the routers maintain IP network connectivity.
<b>Options</b>	<b>timeout <i>seconds</i></b> —Time in seconds before the LDP session is torn down and resigaled. <b>Range:</b> 1 through 65,535 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 LDP Session Protection on page 36</a></li> </ul>

## strict-targeted-hellos

---

<b>Syntax</b>	strict-targeted-hellos;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Prevent LDP sessions from being established with remote neighbors that have not been specifically configured. LDP peers will not respond to targeted hellos coming from a source that is not one of the configured remote neighbors.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Enabling Strict Targeted Hello Messages for LDP on page 10</a></li> </ul>

## targeted-hello

---

<b>Syntax</b>	targeted-hello { hello-interval <i>seconds</i> ; hold-time <i>seconds</i> ; }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.5. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the LDP timer and LDP hold time for targeted hellos.
<b>Options</b>	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>• <a href="#">Configuring the LDP Timer for Hello Messages on page 25</a></li><li>• <a href="#">Configuring the Delay Before LDP Neighbors Are Considered Down on page 26</a></li></ul>

## traceoptions (Protocols LDP)

<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> &lt;flag-modifier&gt; &lt;disable&gt;; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols <i>ldp</i>],  [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols <i>ldp</i>],  [edit protocols <i>ldp</i>],  [edit routing-instances <i>routing-instance-name</i> protocols <i>ldp</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p><b>match-on address</b> option for the <b>filter</b> flag modifier added in Junos OS Release 10.4.</p> <p><b>nsr-synchronization</b> and <b>p2mp-nsr-synchronization</b> operations for <b>flag</b> statement introduced in Junos OS Release 13.3.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	Specify LDP protocol-level trace options.
<b>Default</b>	The default LDP protocol-level trace options are inherited from the routing protocols <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.
<b>Options</b>	<p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>file <i>filename</i></b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>ldp-log</b>. We recommend that you place LDP tracing output in the file <b>ldp-log</b>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</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>flag <i>flag</i></b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <ul style="list-style-type: none"> <li>• <b>address</b>—Operation of address and address withdrawal messages</li> <li>• <b>binding</b>—Label-binding operations</li> <li>• <b>error</b>—Error conditions</li> <li>• <b>event</b>—Protocol events</li> </ul>

- **initialization**—Operation of initialization messages
- **label**—Operation of label request, label map, label withdrawal, and label release messages
- **notification**—Operation of notification messages
- **nsr-synchronization**— Nonstop active routing synchronization events
- **p2mp-nsr-synchronization**—Point-to-multipoint nonstop active routing synchronization events
- **packets**—Equivalent to setting **address**, **initialization**, **label**, **notification**, and **periodic** flags (see also the **filter** flag modifier)
- **path**—Label-switched path operations
- **periodic**—Operation of hello and keepalive messages
- **route**—Operation of route messages
- **state**—Protocol state transitions

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **disable**—Disable this trace flag.
- **filter**—Filter to apply to this flag. The **filter** flag modifier can be applied only to the **route**, **path**, and **binding** flags. This flag modifier has the following options:
  - **match-on**—Match on argument specified. The **match-on** option has the following suboptions:
    - **address**—Filter based on the source and destination addresses of packets. Available for the **packets** flag option only.
    - **fec**—Filter based on the FEC associated with the traced object.
    - **policy *policy-name***—Specify the filter policy.
  - **receive**—Packets being received.
  - **send**—Packets being transmitted.

**no-world-readable**—(Optional) Prevent all users from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

**world-readable**—(Optional) Enable any user to read the log file.

<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 LDP Protocol Traffic on page 17</a></li> <li>• <i>Network Management Administration Guide for Routing Devices</i></li> </ul>

## track-igp-metric

<b>Syntax</b>	track-igp-metric;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols ldp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp], [edit protocols ldp], [edit routing-instances <i>routing-instance-name</i> protocols ldp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Cause the IGP route metric to be used for the LDP routes instead of the default LDP route metric (the default LDP route metric is 1).
<b>Required Privilege Level</b>	interface—To view this statement in the configuration.
	interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring LDP to Use the IGP Route Metric on page 33</a></li> </ul>

## traffic-statistics (Protocols LDP)

<b>Syntax</b>	<pre>traffic-statistics {     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>;     no-penultimate-hop; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],</p> <p>[edit protocols ldp],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols ldp]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	LDP traffic statistics display the amount of traffic passed through a router for a particular FEC.
<b>Options</b>	<p><b>file <i>filename</i></b>—Name of the file to receive the output of the LDP statistics operation. Enclose the name within quotation marks. All files are placed in the directory <code>/var/log</code>.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of LDP statistics files. When a statistics file named <b><i>ldp-stat</i></b> reaches its maximum size, it is renamed <b><i>ldp-stat.0</i></b>, then <b><i>ldp-stat.1</i></b>, and so on, until the maximum number of LDP statistics files is reached. Then the oldest 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 also must include the <b>size</b> statement to specify the maximum file size.</p> <p><b>interval <i>seconds</i></b>—(Optional) Specify the interval at which the statistics are polled and written to the file.</p> <p><b>Default:</b> 300 seconds (5 minutes)</p> <p><b>no-penultimate-hop</b>—(Optional) Do not collect traffic statistics on the penultimate hop router.</p> <p><b>no-world-readable</b>—(Optional) Prevent all users from reading the log file.</p> <p><b>size <i>size</i></b>—(Optional) Maximum size of each statistics file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a statistics file named <b><i>ldp-stat</i></b> reaches this size, it is renamed <b><i>ldp-stat.0</i></b>. When <b><i>ldp-stat</i></b> again reaches this size, <b><i>ldp-stat.0</i></b> is renamed <b><i>ldp-stat.1</i></b> and <b><i>ldp-stat</i></b> is renamed <b><i>ldp-stat.0</i></b>. This renaming scheme continues until the maximum number of statistics files is reached. Then the oldest statistics file is overwritten.</p> <p><b>Syntax:</b> <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>



**Default:** 1 MB

If you specify a maximum file size, you also must also include the **files** statement to specify the maximum number of files.

**world-readable**—(Optional) Enable log file access for all users.

<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Collecting LDP Statistics on page 15</a></li></ul>
------------------------------	--

## transport-address

<b>Syntax</b>	<code>transport-address (interface   router-id);</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols ldp],          [edit logical-systems <i>logical-system-name</i> protocols ldp interface <i>interface-name</i>],          [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols ldp],          [edit protocols ldp],          [edit protocols ldp interface <i>interface-name</i>],          [edit routing-instances <i>routing-instance-name</i> protocols ldp interface <i>interface-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	<p>Enables you to configure the IP address used to specify the TCP session for the LDP session. Routers must first establish a TCP session between one another before they can establish an LDP session. The TCP session enables the routers to exchange the label advertisements needed for the LDP session. To establish the TCP session, each router must learn the other router's transport address. The transport address is an IP address used to identify the TCP session over which the LDP session will run.</p>
<b>Default</b>	<b>router-id</b>
<b>Options</b>	<p><b>interface</b>—The first IP address on the interface is used as the transport address for any LDP sessions to neighbors that can be reached over that interface. You cannot specify the <b>interface</b> option when there are multiple parallel links to the same LDP neighbor, because the LDP specification requires that the same transport address be advertised on all interfaces to the same neighbor. If LDP detects multiple parallel links to the same neighbor, it disables interfaces to that neighbor one by one until the condition is cleared, either by disconnecting the neighbor on an interface or by specifying the <b>router-id</b> option.</p> <p><b>router-id</b>—The router identifier is used as the transport address. Unless otherwise configured, the router identifier is the loopback address.</p>
<b>Required Privilege Level</b>	<p><b>interface</b>—To view this statement in the configuration.</p> <p><b>interface-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Specifying the Transport Address Used by LDP on page 15</a></li> </ul>

## CHAPTER 3

# Monitoring Commands for LDP

- `clear ldp neighbor`
- `clear ldp session`
- `clear ldp statistics`
- `ping mpls ldp`
- `show ldp database`
- `show ldp fec-filters`
- `show ldp interface`
- `show ldp neighbor`
- `show ldp path`
- `show ldp route`
- `show ldp session`
- `show ldp statistics`
- `show ldp traffic-statistics`
- `traceroute mpls ldp`

## clear ldp neighbor

---

<b>Syntax</b>	<code>clear ldp neighbor</code> <code>&lt;instance <i>instance-name</i>&gt;</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code> <code>&lt;neighbor&gt;</code>
<b>Description</b>	Tear down Label Distribution Protocol (LDP) neighbor connections.
<b>Options</b>	<p><b>none</b>—Tear down connections with all LDP neighbors for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear the LDP session for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor</b>—(Optional) Clear an LDP session for the specified neighbor (IP address) only.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show ldp neighbor on page 102</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear ldp neighbor on page 84</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear ldp neighbor

```
user@host> clear ldp neighbor
```

## clear ldp session

<b>Syntax</b>	clear ldp session <destination> <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Clear Label Distribution Protocol (LDP) sessions.
<b>Options</b>	<p><b>none</b>—Clear LDP sessions for all destinations for all routing instances.</p> <p><b>destination</b>—(Optional) Clear an LDP session for the specified destination (IP address).</p> <p><b>instance <i>instance-name</i></b>—(Optional) Clear the LDP session for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ldp session on page 110</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear ldp session on page 85</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear ldp session

```
user@host> clear ldp session
```

## clear ldp statistics

---

<b>Syntax</b>	<code>clear ldp statistics</code> <code>&lt;instance <i>instance-name</i>&gt;</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Set all Label Distribution Protocol (LDP) statistics to zero.
<b>Options</b>	<b>none</b> —Set all LDP statistics to zero for all routing instances.  <b>instance <i>instance-name</i></b> —(Optional) Clear the LDP session for the specified routing instance only.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show ldp statistics on page 116</a></li><li>• <a href="#">show ldp traffic-statistics on page 120</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear ldp statistics on page 86</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear ldp statistics

```
user@host> clear ldp statistics
```

## ping mpls ldp

<b>Syntax</b>	<pre>ping mpls ldp <i>fec</i> &lt;count <i>count</i>&gt; &lt;destination <i>address</i>&gt; &lt;detail&gt; &lt;exp <i>forwarding-class</i>&gt; &lt;instance <i>routing-instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;p2mp root-addr <i>ip-address</i> lsp-id <i>identifier</i>&gt; &lt;size <i>bytes</i>&gt; &lt;source <i>source-address</i>&gt; &lt;sweep&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p><b>size</b> and <b>sweep</b> options introduced in Junos OS Release 9.6.</p> <p><b>instance</b> option introduced in Junos OS Release 10.0.</p> <p><b>p2mp</b>, <b>root-address</b>, and <b>lsp-id</b> options introduced in Junos OS Release 11.2.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	<p>Check the operability of MPLS LDP-signaled label-switched path (LSP) connections. Type Ctrl+c to interrupt a <b>ping mpls</b> command.</p>
<b>Options</b>	<p><b>count</b> <i>count</i>—(Optional) Number of ping requests to send. If <b>count</b> is not specified, five ping requests are sent. The range of values is 1 through <b>1,000,000</b>. The default value is <b>5</b>.</p> <p><b>destination</b> <i>address</i>—(Optional) Specify an address other than the default (<b>127.0.0.1/32</b>) for the ping echo requests. The address can be anything within the <b>127/8</b> subnet.</p> <p><b>detail</b>—(Optional) Display detailed information about the echo requests sent and received.</p> <p><b>exp</b> <i>forwarding-class</i>—(Optional) Value of the forwarding class for the MPLS ping packets.</p> <p><b>fec</b>—Ping an LDP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.</p> <p><b>instance</b> <i>routing-instance-name</i>—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.</p> <p><b>logical-system</b> (all   <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on the specified logical system.</p> <p><b>p2mp root-addr</b> <i>ip-address</i> <b>lsp-id</b> <i>identifier</i>—(Optional) Ping the end points of a point-to-multipoint LSP. Enter the IP address of the point-to-multipoint LSP root and the ID number of the point-to-multipoint LSP.</p> <p><b>size</b> <i>bytes</i>—(Optional) Size of the LSP ping request packet (<b>88</b> through <b>65468</b> bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller</p>

than the minimum size, an error message is displayed reminding you of the 88-byte minimum.

**source source-address**—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

You can configure the ping interval for the **ping mpls ldp** command by specifying a new time in seconds using the **lsp-ping-interval** statement at the **[edit protocols ldp oam]** hierarchy level. For more information, see the *MPLS Applications Feature Guide for Routing Devices*.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls ldp fec count on page 88](#)  
[ping mpls ldp p2mp root-addr lsp-id on page 88](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping mpls ldp fec count

```
user@host> ping mpls ldp 10.255.245.222 count 10
!!!xxx...x--- lsping statistics ---10 packets transmitted, 3 packets received,
70% packet loss 4 packets received with error status, not counted as received.
```

### ping mpls ldp p2mp root-addr lsp-id

```
user@host> ping mpls ldp p2mp root-addr 10.1.1.1/32 lsp-id 1 count 1
Request for seq 1, to interface 71, no label stack.
Request for seq 1, to interface 70, label 299786
Reply for seq 1, egress 10.1.1.3, return code: Egress-ok, time: 18.936 ms
  Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
  Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
```



```
Reply for seq 1, egress 10.1.1.4, return code: Egress-ok, time: 18.936 ms
  Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
  Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
Reply for seq 1, egress 10.1.1.5, return code: Egress-ok, time: 18.936 ms
  Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
  Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
```

## show ldp database

---

<b>Syntax</b>	<code>show ldp database</code> <code>&lt;brief   detail   extensive&gt;</code> <code>&lt;inet   l2circuit&gt;</code> <code>&lt;instance <i>instance-name</i>&gt;</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code> <code>&lt;p2mp&gt;</code> <code>&lt;session <i>session</i>&gt;</code> <code>&lt;summary&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. <b>summary</b> option introduced in Junos OS Release 14.2.
<b>Description</b>	Display entries in the LDP database.
<b>Options</b>	<b>none</b> —Display standard information about all entries in the LDP database for all routing instances.  <b>brief   detail   extensive</b> —(Optional) Display the specified level of output.  <b>inet   l2circuit</b> —(Optional) Display only IPv4 or Layer 2 circuit bindings.  <b>instance <i>instance-name</i></b> —(Optional) Display routing instance information for the specified instance only.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.  <b>p2mp</b> —(Optional) Display point-to-multipoint binding information.  <b>session <i>session</i></b> —(Optional) Display database for the specified session only. <b><i>session</i></b> is the destination address of the LDP session.  <b>summary</b> —(Optional)—Display summary output. This option displays the number of labels received and advertised for each LDP session.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ldp database (master) on page 93</a> <a href="#">show ldp database (standby) on page 94</a> <a href="#">show ldp database l2circuit detail on page 94</a> <a href="#">show ldp database l2circuit extensive on page 95</a> <a href="#">show ldp database p2mp (master) on page 95</a> <a href="#">show ldp database p2mp (standby) on page 95</a> <a href="#">show ldp database p2mp (master) on page 96</a> <a href="#">show ldp database p2mp (standby) on page 96</a> <a href="#">show ldp database session on page 96</a> <a href="#">show ldp database (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 97</a>

[show ldp database \(Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 97](#)  
[show ldp database summary on page 98](#)

**Output Fields** [Table 5 on page 91](#) describes the output fields for the **show ldp database** command. Output fields are listed in the approximate order in which they appear.

**Table 5: show ldp database Output Fields**

Field Name	Field Description	Level of Output
<b>Input label database</b>	Label received from the other router.	All levels
<b>Output label database</b>	Label advertised to the other router.	All levels
<i>session-identifier</i>	Session identifier, which includes the local and remote label space identifiers.	All levels
<b>Labels received</b>	Number of labels received from the other router.	All levels
<b>Labels advertised</b>	Number of labels advertised to the other router.	All levels.
<b>Label</b>	Label binding to a route prefix.	All levels

Table 5: show ldp database Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Prefix</b>	<p>Route prefix.</p> <p>It can be one of the following values:</p> <ul style="list-style-type: none"> <li>IP prefix.</li> <li>Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.</li> <li>Layer 2 encapsulation type.</li> </ul> <p>Layer 2 encapsulation types are displayed in the format <b>L2CKT control word status encapsulation-type vc-number</b>, for example, <b>L2CKT CtlfWord FRAME RELAY VC 2</b></p> <ul style="list-style-type: none"> <li><b>control-word-status</b>—Displays whether the use of the control word has been negotiated for this virtual circuit: <ul style="list-style-type: none"> <li><b>NoCtrlWord</b></li> <li><b>CtrlWord</b></li> </ul> </li> <li><b>encapsulation-type</b>—Encapsulation type: <ul style="list-style-type: none"> <li><b>FRAME RELAY</b></li> <li><b>ATM AAL5</b></li> <li><b>ATM CELL</b></li> <li><b>VLAN</b></li> <li><b>ETHERNET</b></li> <li><b>CISCO_HDLC</b></li> <li><b>PPP</b></li> </ul> </li> <li><b>VC number</b>—Virtual circuit number. It can have any numeric value.</li> <li><b>(Stale)</b>—When you display the LDP database for the neighbor of a restarting router, the bindings learned from the restarting neighbor are displayed as (Stale). Stale bindings are deleted if they are not refreshed within the recovery time.</li> </ul>	All levels
<b>MTU</b>	MTU of the Layer 2 circuit. MTU is displayed for all encapsulation types except ATM cell encapsulations.	<b>detail</b>
<b>VCCV Control Channel types</b>	<p>Virtual Circuit Connection Verification (VCCV) control channel types.</p> <ul style="list-style-type: none"> <li><b>MPLS router alert label</b></li> <li><b>MPLS PW label with TTL=1</b></li> </ul>	<b>extensive</b>
<b>VCCV Control Verification types</b>	The only valid VCCV control verification type is <b>LSP ping</b> .	<b>extensive</b>
<b>TDM payload size</b>	Size of the Time Division Multiplex (TDM) payload.	All levels
<b>TDM bitrate</b>	Bit rate for the TDM traffic.	All levels
<b>Requested VLAN ID</b>	(VLANs) VLAN identifier of the Layer 2 circuit.	<b>detail</b>
<b>Cell bundle size</b>	(ATM cell encapsulations) Maximum number of cells that the Layer 2 circuit can receive in a packet.	<b>detail</b>

Table 5: show ldp database Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	State of the label binding: <ul style="list-style-type: none"> <li><b>Active</b>—Label binding has been installed and distributed appropriately. A label binding is almost always in this state.</li> <li><b>New</b>—New label that has not yet been distributed. <ul style="list-style-type: none"> <li><b>MapRcv</b>—Waiting to receive a label mapping message.</li> <li><b>MapSend</b>—Waiting to send a label mapping message.</li> <li><b>RelRcv</b>—Waiting to receive a label release message.</li> <li><b>RelRsnd</b>—Waiting to receive a label release message before resending label mapping message.</li> <li><b>RelSend</b>—Waiting to send a label release message.</li> <li><b>ReqSend</b>—Waiting to send a label request message.</li> <li><b>W/dSend</b>—Waiting to send a label withdrawal message.</li> </ul> </li> </ul>	<b>detail</b>
<b>Age</b>	Time elapsed since the binding was created.	<b>detail</b>

## Sample Output

### show ldp database (master)

```

user@host> show ldp database extensive
Input label database, 10.255.107.232:0--10.255.107.236:0
  Label Prefix
  299840 10.255.107.232/32
          State: Active
          Age: 9:35
          Entropy Label Capability: No
    3     10.255.107.236/32
          State: Active
          Age: 9:35
          Entropy Label Capability: No
  299776 L2CKT CtrlWord VLAN VC 100
          MTU: 1500 Requested VLAN ID: 600 Flow Label T Bit: 1 Flow Label R
          Bit: 1
          State: Active
          Age: 9:35
          Entropy Label Capability: No
          VCCV Control Channel types:
            PWE3 control word
            MPLS router alert label
            MPLS PW label with TTL=1
          VCCV Control Verification types:
            LSP ping
            BFD with PW-ACH-encapsulation for Fault Detection
            BFD with IP/UDP-encapsulation for Fault Detection

Output label database, 10.255.107.232:0--10.255.107.236:0
  Label Prefix
    3     10.255.107.232/32
          State: Active
          Age: 9:35
          Entropy Label Capability: No
  299776 10.255.107.236/32

```

```

State: Active
Age: 9:35
Entropy Label Capability: No

```

### show ldp database (standby)

```
user@host> show ldp database extensive
```

```

Input label database, 10.255.107.236:0--10.255.107.234:0
Label Prefix
299808 10.255.107.230/32
      State: Active
      Age: 1d 2:46:36
      Standby binding state:
      Map messages: 1
      Release messages: 0
Label Prefix
301136 10.255.107.232/32
      State: Active
      Age: 1d 2:46:36
      Standby binding state:
      Map messages: 1
      Release messages: 0
Label Prefix
3      10.255.107.234/32
      State: Active
      Age: 1d 2:46:36
      Standby binding state:
      Map messages: 1
      Release messages: 0
Label Prefix
302480 10.255.107.236/32
      State: Active
      Age: 1d 2:46:36
      Standby binding state:
      Map messages: 1
      Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.234:0
Label Prefix
299904 10.255.107.230/32
      State: Active
      Age: 1d 2:46:36
299936 10.255.107.232/32
      State: Active
      Age: 1d 2:46:36
299872 10.255.107.234/32
      State: Active
      Age: 1d 2:46:36
3      10.255.107.236/32
      State: Active
      Age: 1d 2:46:36
299952 P2MP root-addr 10.255.107.230, lsp-id 16777217
      State: Active
      Age: 1d 2:46:36

```

### show ldp database l2circuit detail

```

user@host> show ldp database l2circuit detail
Input label database, 10.255.245.44:0--10.255.245.45:0
Label Prefix

```

```

100176      L2CKT CtrlWord ATM CELL (VC Mode) VC 100
            Cell bundle size: 80
            State: Active
            Age: 9:48
100256      L2CKT CtrlWord FRAME RELAY VC 101
            MTU: 4470
            State: Active
            Age: 9:48

Output label database, 10.255.245.44:0--10.255.245.45:0
Label      Prefix
100048      L2CKT CtrlWord ATM CELL (VC Mode) VC 100
            Cell bundle size: 80
            State: Active
            Age: 9:48
100112      L2CKT CtrlWord FRAME RELAY VC 101
            MTU: 4470
            State: Active
            Age: 9:48

```

#### show ldp database l2circuit extensive

```

user@host> show ldp database l2circuit extensive
Input label database, 10.255.245.198:0--10.255.245.194:0
Label      Prefix
299872      L2CKT CtrlWord PPP VC 100
            MTU: 4470
            VCCV Control Channel types:
              MPLS router alert label
              MPLS PW label with TTL=1
            VCCV Control Verification types:
              LSP ping
Label      Prefix
            State: Active
            Age: 19:23:08

```

#### show ldp database p2mp (master)

```

user@host> show ldp database p2mp extensive

Input label database, 10.255.107.232:0--10.255.107.236:0
Label      Prefix
569649      P2MP root-addr 10.255.107.232, lsp-id 16777217
            State: Active
            Age: 2d 6:41:46

Output label database, 10.255.107.232:0--10.255.107.236:0

Input label database, 10.255.107.232:0--10.255.107.238:0

Output label database, 10.255.107.232:0--10.255.107.238:0
Label      Prefix
299888      P2MP root-addr 10.255.107.230, lsp-id 16777217
            State: Active
            Age: 2d 6:41:35

```

#### show ldp database p2mp (standby)

```

user@host> show ldp database p2mp extensive

Input label database, 10.255.107.236:0--10.255.107.232:0

```

```

Label      Prefix
299968     P2MP root-addr 10.255.107.230, lsp-id 16777217
           State: Active
           Age: 4d 22:21:57
           Standby binding state:
             Map messages: 1
             Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.232:0
Label      Prefix
3          P2MP root-addr 10.255.107.232, lsp-id 1
           State: Active
           Age: 4d 22:21:57

```

### show ldp database p2mp (master)

```
user@host> show ldp database p2mp extensive
```

```

Input label database, 10.255.107.232:0--10.255.107.236:0
Label      Prefix
569649     P2MP root-addr 10.255.107.232, lsp-id 16777217
           State: Active
           Age: 2d 6:41:46

Output label database, 10.255.107.232:0--10.255.107.236:0

Input label database, 10.255.107.232:0--10.255.107.238:0

Output label database, 10.255.107.232:0--10.255.107.238:0
Label      Prefix
299888     P2MP root-addr 10.255.107.230, lsp-id 16777217
           State: Active
           Age: 2d 6:41:35

```

### show ldp database p2mp (standby)

```
user@host> show ldp database p2mp extensive
```

```

Input label database, 10.255.107.236:0--10.255.107.232:0
Label      Prefix
299968     P2MP root-addr 10.255.107.230, lsp-id 16777217
           State: Active
           Age: 4d 22:21:57
           Standby binding state:
             Map messages: 1
             Release messages: 0

Output label database, 10.255.107.236:0--10.255.107.232:0
Label      Prefix
3          P2MP root-addr 10.255.107.232, lsp-id 1
           State: Active
           Age: 4d 22:21:57

```

### show ldp database session

```

user@host> show ldp database session 10.1.1.195
Input label database, 10.0.0.194:0--10.1.1.195:0
Label      Prefix
100002     10.255.245.197/32
100003     10.255.245.196/32
100004     10.0.0.194/32

```



```

      3      10.1.1.195/32
100000      L2CKT NoCtrlWord FRAME RELAY VC 1
100001      L2CKT CtrlWord FRAME RELAY VC 2
Output label database, 10.0.0.194:0--10.1.1.195:0
  Label      Prefix
100003      10.255.245.197/32
100004      10.1.1.195/32
100002      10.255.245.196/32
      3      10.0.0.194/32
100000      L2CKT CtrlWord FRAME RELAY VC 2
100001      L2CKT NoCtrlWord FRAME RELAY VC 1

```

#### show ldp database (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show ldp database
Input label database, 1.1.1.2:0--1.1.1.3:0
  Label      Prefix
299808      1.1.1.2/32
      3      1.1.1.3/32
299792      1.1.1.6/32
299776      10.255.2.227/32
299840      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
299824      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11

Output label database, 1.1.1.2:0--1.1.1.3:0
  Label      Prefix
      3      1.1.1.2/32
299776      1.1.1.3/32
299808      1.1.1.6/32
299792      10.255.2.227/32

Input label database, 1.1.1.2:0--1.1.1.6:0
  Label      Prefix
299856      1.1.1.2/32
299792      1.1.1.3/32
      3      1.1.1.6/32
299776      10.255.2.227/32
299888      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
299808      P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
299824      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
299840      P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
299872      P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7

Output label database, 1.1.1.2:0--1.1.1.6:0
  Label      Prefix
      3      1.1.1.2/32
299776      1.1.1.3/32
299808      1.1.1.6/32
299792      10.255.2.227/32

```

#### show ldp database (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show ldp database
Input label database, 10.255.2.227:0--1.1.1.3:0
  Label      Prefix
299808      1.1.1.2/32
      3      1.1.1.3/32
299792      1.1.1.6/32
299776      10.255.2.227/32

Output label database, 10.255.2.227:0--1.1.1.3:0
  Label      Prefix

```

```

299856      1.1.1.2/32
299776      1.1.1.3/32
299792      1.1.1.6/32
3           10.255.2.227/32

```

Input label database, 10.255.2.227:0--1.1.1.6:0

```

Label      Prefix
299856      1.1.1.2/32
299776      1.1.1.3/32
3           1.1.1.6/32
299776      10.255.2.227/32

```

Output label database, 10.255.2.227:0--1.1.1.6:0

```

Label      Prefix
299856      1.1.1.2/32
299776      1.1.1.3/32
299792      1.1.1.6/32
3           10.255.2.227/32
299888      P2MP root-addr 1.1.1.2, grp: 232.2.2.2, src: 1.2.7.7
299808      P2MP root-addr 1.1.1.2, grp: 232.1.1.1, src: 192.168.219.11
299824      P2MP root-addr 1.1.1.2, grp: 232.1.1.2, src: 192.168.219.11
299840      P2MP root-addr 1.1.1.2, grp: 232.1.1.3, src: 192.168.219.11
299872      P2MP root-addr 1.1.1.2, grp: ff3e::1:2, src: abcd::1:2:7:7

```

#### show ldp database summary

```
user@host> show ldp database summary
```

Session ID	Labels received	Labels advertised
10.255.0.1:0--10.255.0.2:0	4	4
10.255.0.1:0--10.255.0.3:0	4	4

## show ldp fec-filters

<b>Syntax</b>	show ldp fec-filters <fec> <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Display information about configured Label Distribution Protocol (LDP) forwarding equivalence class (FEC) filters.
<b>Options</b>	<p><b>fec</b>—(Optional) Display FEC filter information for the specified FEC.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display FEC filter information for the specified instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ldp fec-filters on page 99</a>
<b>Output Fields</b>	<a href="#">Table 6 on page 99</a> lists the output fields for the <b>show ldp fec-filters</b> command. Output fields are listed in the approximate order in which they appear.

**Table 6: show ldp fec-filters Output Fields**

Field Name	Field Description
Ingress	Names of the FEC filters on the ingress routers.
Transit	Names of the FEC filters on the transit routers.

## Sample Output

### show ldp fec-filters

```
user@host> show ldp fec-filters 10/8
10.22.1.2/32
  Ingress: f1-10.22.1.2/32 (index: 3)
  Transit: (null) (index: 0)
```

## show ldp interface

<b>Syntax</b>	<pre>show ldp interface &lt;brief   detail   extensive&gt; &lt;interface-name&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	Display the status of Label Distribution Protocol (LDP)-enabled interfaces.
<b>Options</b>	<p><b>none</b>—Display standard status information about all LDP-enabled interface for all routing instances.</p> <p><b>interface-name</b>—(Optional) Display information for the specified interface.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>instance instance-name</b>—(Optional) Display information for the specified routing instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ldp interface extensive on page 101</a>
<b>Output Fields</b>	<p><a href="#">Table 7 on page 100</a> describes the output fields for the <b>show ldp interface</b> command. Output fields are listed in the approximate order in which they appear.</p>

Table 7: show ldp interface Output Fields

Field Name	Field Description	Level of Output
<b>Interface</b>	Interface name.	All levels
<b>Label space ID</b>	Label space identifier that the router is advertising on the interface.	All levels
<b>Nbr count</b>	Number of neighbors on the interface.	All levels
<b>Next hello</b>	How long until the next hello packet is sent on this interface, in seconds.	All levels
<b>Hello interval</b>	One-third of the negotiated hold time (in seconds). If the user-configured value for the hello interval is smaller than the computed value, the user-configured value is used.	<b>detail</b> <b>extensive</b>
<b>Hold time</b>	Configured hold time, in seconds.	<b>detail</b> <b>extensive</b>

Table 7: show ldp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transport address	Address to which the neighbor wants the local route to establish the LDP session.	extensive
Local hello interval	Locally configured hello interval.	extensive

## Sample Output

### show ldp interface extensive

```
user@host> show ldp interface extensive
Interface          Label space ID      Nbr count  Next hello
fe-0/0/3.0         10.255.245.6:0      2           0
Hello interval: 1, Hold time: 15, Transport address: 10.255.245.6
Local hello interval: 2, Index: 69
```

## show ldp neighbor

<b>Syntax</b>	<pre>show ldp neighbor &lt;brief   detail   extensive&gt; &lt;auto-targeted&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;neighbor-address&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p><b>neighbor-address</b> option added in Junos OS Release 8.5.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p> <p><b>auto-targeted</b> option added in Junos OS Release 14.2.</p>
<b>Description</b>	Display Label Distribution Protocol (LDP) neighbor information.
<b>Options</b>	<p><b>none</b>—Display standard information about LDP neighbors for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>auto-targeted</b>—(Optional) Display information about LDP neighbors that are automatically targeted using the loopback addresses.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified routing instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>neighbor-address</b>—(Optional) Display information about the specified LDP neighbor.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear ldp neighbor on page 84</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show ldp neighbor extensive on page 103</a></p> <p><a href="#">show ldp neighbor auto-targeted extensive on page 103</a></p>
<b>Output Fields</b>	<p><a href="#">Table 8 on page 102</a> describes the output fields for the <b>show ldp neighbor</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 8: show ldp neighbor Output Fields**

Field Name	Field Description	Level of Output
<b>Address</b>	IP address of the neighbor.	All levels
<b>Interface</b>	Interface over which the neighbor was discovered.	All levels
<b>Label space ID</b>	Label space identifier advertised by the neighbor.	All levels

Table 8: show ldp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Hold time	Remaining hold time before the neighbor expires, in seconds.	All levels
Transport address	Address to which the neighbor wants the local route to establish the LDP session.	detail
Configuration sequence	Counter that increments whenever the neighbor changes its configuration.	detail
Up for	Length of time the LDP neighbor has been in operation.	detail extensive
Reference count	Reference count for the LDP neighbor.	extensive
Hold time	Displays the neighbor's hold time. The hold time is the proposed hold times for the local and peer routers.	extensive
Proposed local/peer	Hold time value proposed by the local router and the peer router.	extensive

## Sample Output

### show ldp neighbor extensive

```

user@host> show ldp neighbor extensive
Address          Interface      Label space ID      Hold Time
192.168.37.23    so-1/0/0.0    10.255.245.5:0      44
Transport address: 10.255.245.5, Configuration sequence: 6
Up for 00:03:37
Reference count: 1
Hold time: 45, Proposed local/peer: 15/45

```

### show ldp neighbor auto-targeted extensive

```

user@host> show ldp neighbor auto-targeted extensive
Address          Interface      Label space ID      Hold time
10.255.107.236   lo0.0         10.255.107.236:0    41
Transport address: 10.255.107.236, Configuration sequence: 14
Up for 00:10:53
Reference count: 2
Hold time: 45, Proposed local/peer: 45/45
Hello interval: 15
Hello flags: targeted
Neighbor types: Auto-targeted

```

## show ldp path

<b>Syntax</b>	<pre>show ldp path &lt;brief   detail   extensive&gt; &lt;destination&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	Display Label Distribution Protocol (LDP) label-switched paths (LSPs).
<b>Options</b>	<p><b>none</b>—Display standard information about all LDP LSPs for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>destination</b>—(Optional) Restrict the output to entries that match the specified destination prefix.</p> <p><b>instance instance-name</b>—(Optional) Display information for the specified routing instance only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ldp path extensive on page 105</a>
<b>Output Fields</b>	<a href="#">Table 9 on page 104</a> describes the output fields for the <b>show ldp path</b> command. Output fields are listed in the approximate order in which they appear.

**Table 9: show ldp path Output Fields**

Field Name	Field Description
<b>Output Session (label)</b>	Session ID and labels that this system has sent using LDP. These correspond to MPLS packets received.
<b>Input Session (label)</b>	Session ID and labels that this system has received using LDP. These correspond to MPLS packets transmitted.
<b>route</b>	MPLS route.
<b>Attached route</b>	Route corresponding to the LSP.
<b>Ingress route</b>	The router acts as the ingress for the LSP.
<b>Reference count</b>	Reference count for the LDP neighbor.



Table 9: show ldp path Output Fields (*continued*)

Field Name	Field Description
<b>Transit route</b>	Names of the forwarding equivalence class (FEC) filters on the transit routers.
<b>Global label</b>	MPLS label that is used globally.

## Sample Output

### show ldp path extensive

```

user@host> show ldp path extensive
Output Session (label)      Input Session (label)
10.255.14.220:0(3)         ( )
  Attached route: 10.255.14.221/32
  Reference count: 3, Global label: 3
10.255.14.220:0(100000)     10.255.14.220:0(3)
  Attached route: 10.255.14.220/32, Ingress route
  Reference count: 2, Transit route, Global label: 100000
10.255.14.220:0(100001)     10.255.14.220:0(100001)
  Attached route: 10.255.14.214/32, Ingress route
  Reference count: 2, Transit route, Global label: 100001

```

## show ldp route

<b>Syntax</b>	<pre>show ldp route &lt;brief   detail   extensive&gt; &lt;destination&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	Display the entries in the Label Distribution Protocol (LDP) internal topology table. The internal topology table contains routes from inet.0 and inet.3 and is used when binding a label to a forwarding equivalence class (FEC).
<b>Options</b>	<p><b>none</b>—Display standard information about all entries in the LDP internal topology table for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>destination</b>—(Optional) Restrict the output to entries that are longer than the specified destination prefix and prefix length.</p> <p><b>instance instance-name</b>—(Optional) Display entries for the specified routing instance only.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show ldp route detail on page 108</a></p> <p><a href="#">show ldp route extensive on page 108</a></p>
<b>Output Fields</b>	<a href="#">Table 10 on page 106</a> describes the output fields for the <b>show ldp route</b> command. Output fields are listed in the approximate order in which they appear.

**Table 10: show ldp route Output Fields**

Field Name	Field Description
<b>Destination</b>	Destination prefix.
<b>Next-hop intf/lsp/table</b>	Interface that is the next hop to the destination prefix.
<b>Next-hop address</b>	IP address of the next hop.
<b>Session ID</b>	LDP session ID.

Table 10: show ldp route Output Fields (*continued*)

Field Name	Field Description
Route flags	Information about the route. For example, the <b>Ingress TTL propagate</b> flag indicates that the time-to-live (TTL) value is being propagated with the route.
Bound to outgoing label	The route has been bound to LSPs with the label being distributed for that LSP.
Topology entry	The topology that the route is bound to.
Ingress route status	Status of the ingress route. For example, it could be <b>Active</b> or <b>Inactive</b> .
Last modified	The length of time since the ingress route status last changed.

## Sample Output

### show ldap route detail

```

user@host> show ldap route 10.255.8.5 detail
Destination      Next-hop intf/lsp      Next-hop address
10.255.8.5/32     f1
  Session ID 10.255.170.84:0--10.255.170.92:0
                    fe-0/0/0.0      192.168.100.2
  Session ID 10.255.170.84:0--10.255.8.5:0
                    so-0/2/1.0
  Session ID 10.255.170.84:0--10.255.8.5:0
                    so-0/2/2.0
  Session ID 10.255.170.84:0--10.255.8.3:0
  Bound to outgoing label 299776, Topology entry: 0x8c38a80
  BFD dest addr   BFD state LSP-ping Next-hop addr Next-hop intf/lsp
127.0.0.64       up        up        192.168.100.2 fe-0/0/0.0
127.0.1.64       up        up        so-0/2/1.0
127.0.2.64       up        up        so-0/2/2.0
127.0.3.64       up        up        f1
.....

```

### show ldap route extensive

```

user@host> show ldap route extensive

Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.0/30      ge-1/2/0.18            10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.4/30      ge-1/2/0.18            10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.8/30      ge-1/2/1.21            10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.12/30     ge-1/2/1.21            10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.16/30     ge-1/2/0.18            10.0.0.17
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.18/32     ge-1/2/0.18            10.0.0.17
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.20/30     ge-1/2/1.21            10.0.0.22
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
10.0.0.21/32     ge-1/2/1.21            10.0.0.22
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
192.168.0.1/32   ge-1/2/0.18            10.0.0.17
  Session ID 192.168.0.6:0--192.168.0.5:0
  Route flags: None
Destination      Next-hop intf/lsp/table Next-hop address
192.168.0.2/32   ge-1/2/1.21            10.0.0.22
  Session ID 192.168.0.6:0--192.168.0.4:0

```

```

                                ge-1/2/0.18                10.0.0.17
    Session ID 192.168.0.6:0--192.168.0.5:0
    Route flags: None
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.3/32   ge-1/2/1.21                10.0.0.22
    Session ID 192.168.0.6:0--192.168.0.4:0
    Route flags: None
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.4/32   ge-1/2/1.21                10.0.0.22
    Session ID 192.168.0.6:0--192.168.0.4:0
    Bound to outgoing label 299808, Topology entry: 0x92a483c
    Ingress route status: Active, Last modified: 00:01:19 ago
    Route flags: Ingress TTL propagate, Transit TTL propagate
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.5/32   ge-1/2/0.18                10.0.0.17
    Session ID 192.168.0.6:0--192.168.0.5:0
    Bound to outgoing label 299792, Topology entry: 0x92a47f8
    Ingress route status: Active, Last modified: 00:01:19 ago
    Route flags: Ingress TTL propagate, Transit TTL propagate
Destination      Next-hop intf/lsp/table      Next-hop address
192.168.0.6/32   lo0.6
    Bound to outgoing label 3, Topology entry: 0x92a4a5c
    Ingress route status: Inactive
    Route type: Egress route
    Route flags: None
Destination      Next-hop intf/lsp/table      Next-hop address
10.10.20.1/32    fe-1/0/0.0                192.168.199.37
                                LSP LDP->10.255.107.230

```

## show ldp session

<b>Syntax</b>	<pre>show ldp session &lt;brief   detail   extensive&gt; &lt;auto-targeted&gt; &lt;destination&gt; &lt;instance instance-name&gt; &lt;logical-system (all   logical-system-name)&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p> <p><b>auto-targeted</b> option added in Junos OS Release 14.2.</p>
<b>Description</b>	Display information about Label Distribution Protocol (LDP) sessions.
<b>Options</b>	<p><b>none</b>—Display standard information about all LDP sessions for all routing instances.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>auto-targeted</b>—(Optional) Display information about LDP sessions that are automatically targeted using loopback addresses.</p> <p><b>destination</b>—(Optional) Restrict LDP session display to the specified address.</p> <p><b>instance instance-name</b>—(Optional) Display routing instance information for the specified instance. If <b>instance-name</b> is omitted, information is displayed for the master instance.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear ldp session on page 85</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show ldp session brief on page 114</a></p> <p><a href="#">show ldp session detail on page 114</a></p> <p><a href="#">show ldp session extensive on page 114</a></p> <p><a href="#">show ldp session auto-targeted detail on page 115</a></p>
<b>Output Fields</b>	Table 11 on page 110 describes the output fields for the <b>show ldp session</b> command. Output fields are listed in the approximate order in which they appear.

Table 11: show ldp session Output Fields

Field Name	Field Description	Level of Output
Address	Transport address of the session.	any
State	State of the session: <b>Nonexistent</b> , <b>Connecting</b> , <b>Initialized</b> , <b>OpenRec</b> , <b>OpenSent</b> , <b>Operational</b> , or <b>Closing</b> . The states correspond to the state diagram specified in Internet Draft LDP Specification draft-ietf-mpls-rfc3036bis-01.txt.	any

Table 11: show ldp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Connection	TCP connection state: <b>Closed</b> , <b>Opening</b> , or <b>Open</b> .	any
Hold time	Time remaining until the session will be closed, in seconds.	any
Session ID	LDP identifiers of the peers of this session.	detail extensive
Next keepalive	Time until next keepalive is sent, in seconds.	detail extensive
Active	Whether the local router is playing the active role in the session and during session establishment.	detail extensive
Passive	Whether the local router is playing the passive role in the session and during session establishment.	detail extensive
Maximum PDU	Maximum protocol data unit (PDU) size (packet size) for the session.	detail extensive
Hold time	Time remaining until the session will be closed, in seconds. This value corresponds to the one configured using the <b>keepalive-timeout</b> statement configured at the <b>[edit protocols ldp]</b> hierarchy level.	detail extensive
Neighbor count	Number of neighbors that are contributing to the session.	detail extensive
Neighbor types	Category of LDP session: <b>discovered</b> or <b>auto-targeted</b> .	any
Keepalive interval	Keepalive interval, in seconds.	detail extensive
Connect retry interval	TCP connection retry interval, in seconds.	detail extensive
Local address	Local transport address.	detail extensive
Remote address	Remote transport address.	detail extensive
Up for	Time that this session has been up.	detail extensive
Last down	Time since the session last went down.	detail extensive

Table 11: show ldp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Reason	Reason the session went down: <ul style="list-style-type: none"> <li>Aborted graceful restart</li> <li>Authentication key was changed</li> <li>Bad type length value (TLV)</li> <li>Bad protocol data unit (PDU) packets</li> <li>Command-line interface (CLI) command</li> <li>Connect time expired</li> <li>Connection error</li> <li>Connection reset</li> <li>Error during initialization</li> <li>Hold time expired</li> <li>No adjacency or all adjacencies down</li> <li>Notification received</li> <li>Received notification from peer</li> <li>Unexpected End of File (EOF)</li> <li>Unknown reason</li> </ul>	detail extensive
Number of session flaps	Number of times the session changes from up to down.	detail extensive
Restarting	LDP is in the process of gracefully restarting.	detail extensive
Capabilities advertised	LDP capabilities advertised to a peer.	detail extensive
Capabilities received	LDP capabilities received from a peer.	detail extensive
Protection	Information about the status of MPLS LDP session protection.	detail extensive
restart complete in <i>nnn msec</i>	Amount of time (in milliseconds) remaining until graceful restart is declared complete.	detail extensive
Local	Information about graceful restart for the local end of an LDP session. Graceful restart and helper mode are independent. <ul style="list-style-type: none"> <li><b>Restart</b>—Status of the graceful restart feature at the local end of the LDP session: <b>enabled</b> or <b>disabled</b>.</li> <li><b>Helper mode</b>—Status of the helper mode feature at the local end of the LDP session: <b>enabled</b> or <b>disabled</b>. When this feature is enabled, the local end of the LDP session can help the restarting router with its LDP restart procedures.</li> <li><b>Reconnect time</b>—Amount of time to wait from when a restart is initiated until the router can exchange LDP messages with its neighbors. The default is <b>60000 msec</b> and is not configurable. (<b>Reconnect timeout</b> refers to "FT Reconnect timeout" in draft-ietf-mpls-ldp-restart-06, <i>Internet Draft Graceful Restart Mechanism for LDP</i>.)</li> </ul>	detail extensive



Table 11: show ldp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Remote</b>	Information about graceful restart at the remote end of an LDP session. Graceful restart and helper mode are independent. <ul style="list-style-type: none"> <li>• <b>Restart</b>—Status of the graceful restart feature at the remote end of the LDP session: <b>enabled</b> or <b>disabled</b>.</li> <li>• <b>Helper mode</b>—Status of the helper mode feature at the remote end of the LDP session: <b>enabled</b> or <b>disabled</b>. When this feature is enabled, the remote end of the LDP session can help the restarting router with its LDP restart procedures.</li> <li>• <b>Reconnect time</b>—Amount of time in milliseconds from when a restart is initiated until the remote router can exchange LDP messages with its neighbors.</li> </ul>	<b>detail extensive</b>
<b>Local maximum recovery time</b>	Amount of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds).	<b>detail extensive</b>
<b>Next-hop addresses received</b>	Next-hop addresses received on the session.	<b>detail extensive</b>
<b>Queue depth</b>	Number of messages that are queued for sending to the peers in the group.	<b>extensive</b>
<b>Message type</b>	Type of message being sent: <ul style="list-style-type: none"> <li>• <b>Initialization</b>—Session initialization negotiation messages sent by an LSR to an LDP peer when the transport connection is established.</li> <li>• <b>Keepalive</b>—Keepalive timer messages sent by an LSR to an LDP peer to keep the session active when there is no information or PDU exchanged between them.</li> <li>• <b>Notification</b>—Notification messages (such as state of the LDP session) or error information (such as bad PDU length) sent by an LSR to an LDP peer.</li> <li>• <b>Address</b>—Message sent by an LSR to an LDP peer to advertise interface addresses.</li> <li>• <b>Address withdraw</b>—Message sent by an LSR to an LDP peer to withdraw a previously advertised interface address.</li> <li>• <b>Label mapping</b>—Message sent by an LSR to an LDP peer to advertise label mapping for a forwarding equivalence class (FEC).</li> <li>• <b>Label request</b>—Message sent by an LSR to an LDP peer to request a label mapping for an FEC.</li> <li>• <b>Label withdraw</b>—Message sent by an LSR to an LDP peer to withdraw a previously advertised FEC-label mapping.</li> <li>• <b>Label release</b>—Message sent by an LSR to an LDP peer to notify the peer that a specific FEC-label mapping has been released.</li> <li>• <b>Label abort</b>—Message sent by an LSR to an LDP peer to abort a label request message.</li> <li>• <b>Total</b>—Messages sent and received during the lifetime of the session.</li> <li>• <b>Last 5 seconds</b>—Messages sent and received during the current session.</li> </ul>	<b>extensive</b>

## Sample Output

### show ldp session brief

```
user@host> show ldp session brief
  Address           State           Connection      Hold time
10.255.72.160       Operational     Open            21
10.255.72.164       Operational     Open            20
10.255.72.172       Operational     Open            21
```

### show ldp session detail

```
user@host> show ldp session detail
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 27
Session ID: 192.168.0.2:0--192.168.0.3:0
Next keepalive in 7 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.0.2, Remote address: 192.168.0.3
Up for 00:00:02
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
  10.0.0.5
  10.0.0.33
```

### show ldp session extensive

```
user@host> show ldp session extensive
Address: 192.168.0.3, State: Operational, Connection: Open, Hold time: 22
Session ID: 192.168.0.2:0--192.168.0.3:0
Next keepalive in 2 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.0.2, Remote address: 192.168.0.3
Up for 00:05:37
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Local - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Remote - Restart: enabled, Helper mode: enabled, Reconnect time: 60000
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
Nonstop routing state: Not in sync
Next-hop addresses received:
  10.0.0.5
  10.0.0.33
```

Queue depth: 0

Message type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Initialization	1	1	0	0
Keepalive	33	33	1	1
Notification	0	0	0	0
Address	1	1	0	0
Address withdraw	0	0	0	0
Label mapping	7	5	0	0
Label request	0	0	0	0
Label withdraw	3	1	0	0
Label release	1	3	0	0
Label abort	0	0	0	0

#### show ldp session auto-targeted detail

```

user@host> show ldp session auto-generated detail
Address: 192.168.1.5, State: Operational, Connection: Open, Hold time: 25
  Session ID: 192.168.1.1:0--192.168.1.5:0
  Next keepalive in 5 seconds
  Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
  Neighbor types: discovered, Auto-targeted
                    ^^^^^^^^^^^^^^^^^
  Keepalive interval: 10, Connect retry interval: 1
  Local address: 192.168.1.1, Remote address: 192.168.1.5
  Up for 00:00:34
  Capabilities advertised: none
  Capabilities received: none
  Protection: disabled
  Local - Restart: disabled, Helper mode: enabled
  Remote - Restart: disabled, Helper mode: enabled
  Local maximum neighbor reconnect time: 120000 msec
  Local maximum neighbor recovery time: 240000 msec
  Local Label Advertisement mode: Downstream unsolicited
  Remote Label Advertisement mode: Downstream unsolicited
  Negotiated Label Advertisement mode: Downstream unsolicited
  Nonstop routing state: Not in sync
  Next-hop addresses received:
    192.168.1.2
    192.168.1.3

```

## show ldp statistics

<b>Syntax</b>	show ldp statistics <instance <i>instance-name</i> > <logical-system (all   <i>logical-system-name</i> )>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Display Label Distribution Protocol (LDP) statistics.
<b>Options</b>	<p><b>none</b>—Display LDP statistics for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">clear ldp statistics on page 86</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ldp statistics on page 119</a>
<b>Output Fields</b>	<a href="#">Table 12 on page 116</a> lists the output fields for the <b>show ldp statistics</b> command. Output fields are listed in the approximate order in which they appear.

Table 12: show ldp statistics Output Fields

Field Name	Field Description
Total Sent, Received	Total number of each message type sent and received.
Last 5 seconds Sent, Received	Number of each message type sent and received in the last 5 seconds.

Table 12: show ldp statistics Output Fields (*continued*)

Field Name	Field Description
<b>Message type</b>	<p>LDP message types:</p> <ul style="list-style-type: none"> <li>• <b>Hello</b>—Messages that enable LDP nodes to discover one another and to detect the failure of a neighbor or of the link to the neighbor.</li> <li>• <b>Initialization</b>—Messages that indicate an LDP session has started.</li> <li>• <b>Keepalive</b>—Messages that ensure that the keepalive timeout is not exceeded.</li> <li>• <b>Notification</b>—Advisory information and signal error information.</li> <li>• <b>Address</b>—Messages with address information.</li> <li>• <b>Address withdrawal</b>—Messages regarding address withdrawal.</li> <li>• <b>Label mapping</b>—Messages with label mapping information.</li> <li>• <b>Label request</b>—Request for a label mapping from a neighboring router.</li> <li>• <b>Label withdrawal</b>—Withdrawal message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use.</li> <li>• <b>Label release</b>—Message sent by the downstream LSR to recall a label that it previously mapped. If an LSR that has received a label mapping subsequently determines that it no longer needs that label, it can send a label release message that frees the label for use.</li> <li>• <b>Label abort</b>—Messages about label interruptions.</li> <li>• <b>All UDP</b>—All hello messages sent by LSRs to the well-known UDP port, 646.</li> <li>• <b>All TCP</b>—All LDP session messages.</li> </ul>

Table 12: show ldp statistics Output Fields (*continued*)

Field Name	Field Description
<b>Event type</b>	<p>LDP events and errors:</p> <ul style="list-style-type: none"> <li>• <b>Sessions opened</b>—Number of LDP sessions that have been opened.</li> <li>• <b>Sessions closed</b>—Number of LDP sessions that have been closed.</li> <li>• <b>Topology changes</b>—Number of changes to the known LDP topology.</li> <li>• <b>No interface</b>—Number of missing interface address messages. When a new LDP session is initialized and before sending label lapping or label request messages, the LSR advertises its interface addresses with one or more address messages.</li> <li>• <b>No session</b>—Number of missing session messages. Session messages are used to establish, maintain, and terminate sessions between LDP peers.</li> <li>• <b>No adjacency</b>—The exchange of hello adjacency messages results in the creation of an adjacency. The LDP identifier, together with the sender's LDP identifier in the PDU header, enables the receiver to match the initialization message with one of its hello adjacencies. If there is no matching hello adjacency, the LSR sends a session the initialization message is rejected.</li> <li>• <b>Unknown version</b>—The LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment.</li> <li>• <b>Malformed PDU</b>—An LDP PDU received on a TCP connection for an LDP session is malformed if the LDP identifier in the PDU header is unknown to the receiver, or if it is known but is not the LDP identifier associated by the receiver with the LDP peer for this LDP session.  An LDP PDU is considered to be malformed if the LDP protocol version is not supported by the receiver, or it is supported but is not the version negotiated for the session during session establishment.  An LDP PDU is considered malformed if the PDU length field is too small (less than 14) or too large (greater than maximum PDU length).</li> <li>• <b>Malformed message</b>—Malformed LDP messages that are part of the LDP discovery mechanism are handled by silently discarding them.  An LDP message is malformed if the message type is unknown. If the message type is less than 0x8000 (high order bit = 0), it is an error signaled by the unknown message type status code.  An LDP message is considered to be malformed if the message length is too large, meaning that the message extends beyond the end of the containing LDP PDU.  The LDP message is considered to be malformed if the message length is too small, meaning that it is smaller than the smallest possible value component.  The LDP message is considered to be malformed if the message is missing one or more mandatory parameters.</li> <li>• <b>Unknown message type</b>—If the message type is less than 0x8000 (high order bit = 0) or greater than or equal to 0x8000 (high order bit = 1) it is considered to be an unknown message.</li> <li>• <b>Inappropriate message</b>—The message is not of the type that the receiver expects to receive.</li> <li>• <b>Malformed TLV</b>—The TLV Length is too large or the receiver cannot decode the TLV value. This can indicate an issue in either the sending or receiving LSR.</li> <li>• <b>Bad TLV value</b>—The TLV Length is too large.</li> <li>• <b>Missing TLV</b>—The TLV is missing one or more mandatory parameters.</li> <li>• <b>PDU too large</b>—The PDF is greater than the maximum PDU length. Section "Initialization Message" in RFC 5036 describes how the maximum PDU length for a session is determined.</li> </ul>
<b>Total</b>	Total number of each event or error.
<b>Last 5 seconds</b>	Number of each event or error in the last 5 seconds.

## Sample Output

### show ldp statistics

```
user@host> show ldp statistics
```

Message type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Hello	265	263	2	2
Initialization	2	2	0	0
Keepalive	112	111	1	0
Notification	0	0	0	0
Address	2	2	0	0
Address withdraw	0	0	0	0
Label mapping	7	6	0	0
Label request	0	0	0	0
Label withdraw	2	0	0	0
Label release	0	2	0	0
Label abort	0	0	0	0
All UDP	265	263	2	2
All TCP	123	121	1	0

Event type	Total	Last 5 seconds	
Sessions opened	2		0
Sessions closed	0		0
Topology changes	11		0
No interface	0		0
No session	0		0
No adjacency	0		0
Unknown version	0		0
Malformed PDU	0		0
Malformed message	0		0
Unknown message type	0		0
Inappropriate message	0		0
Malformed TLV	0		0
Bad TLV value	0		0
Missing TLV	0		0
PDU too large	0		0

## show ldp traffic-statistics


<b>Syntax</b>	<pre>show ldp traffic-statistics &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;p2mp&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p><b>p2mp</b> option added in Junos OS Release 11.2.</p> <p>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p>
<b>Description</b>	Display Label Distribution Protocol (LDP) traffic statistics.
<div>  <b>NOTE:</b> If nonstop active routing features is configured, <b>show ldp traffic-statistics</b> command is not supported on backup Routing Engines. </div>	
<b>Options</b>	<p><b>none</b>—Display LDP traffic statistics for all routing instances.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display LDP traffic statistics for the specified routing instance only.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>p2mp</b>—(Optional) Display only the data traffic statistics for a point-to-multipoint LSP.</p>
<b>Additional Information</b>	To collect output from this command on a periodic basis, configure the <a href="#">traffic-statistics</a> statement for the LDP protocol. For more information, see the <i>Junos MPLS Applications Configuration Guide</i> .
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear ldp statistics on page 86</a></li> <li>• <i>Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain</i></li> <li>• <i>Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs</i></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show ldp traffic-statistics on page 121</a></p> <p><a href="#">show ldp traffic-statistics p2mp on page 122</a></p> <p><a href="#">show ldp traffic-statistics p2mp (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 122</a></p> <p><a href="#">show ldp traffic-statistics p2mp (Multipoint LDP with Multicast-Only Fast Reroute) on page 122</a></p>
<b>Output Fields</b>	<p><a href="#">Table 13 on page 121</a> lists the output fields for the <b>show ldp traffic-statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>



Table 13: show ldp traffic-statistics Output Fields

Field Name	Field Description
<b>Message type</b>	LDP message types.
<b>FEC</b>	Forwarding equivalence class (FEC) for which LDP traffic statistics are collected.  For P2MP LSPs, FEC appears as a combination of root address and the LSP ID ( <b>root_addr:lsp_id</b> ).  For M-LDP P2MP LSPs, FEC appears as a combination of root address multicast source address, and multicast group address ( <b>root_addr:lsp_id/grp,src</b> ).
<b>Type</b>	Type of traffic originating from a router, either <b>Ingress</b> (originating from this router) or <b>Transit</b> (forwarded through this router).
<b>Packets</b>	Number of packets passed by the FEC since its LSP came up.
<b>Bytes</b>	Number of bytes of data passed by the FEC since its LSP came up.
<b>Shared</b>	Whether a label is shared by prefixes: <b>Yes</b> or <b>No</b> . A <b>Yes</b> value indicates that several prefixes are bound to the same label (for example, when several prefixes are advertised with an egress policy). The LDP traffic statistics for this case apply to all the prefixes and should be treated as such.
<b>Nextthop</b>	The next hop address for P2MP LSPs. (This is the downstream LDP Session ID.)
<b>Label</b>	For multipoint LDP with multicast-only fast reroute (MoFRR), the multipoint LDP node selects two separate upstream peers and sends two separate labels, one to each upstream peer. The same algorithm described in RFC 6388 is used to select the primary upstream path. The backup upstream path selection again uses the same algorithm but excludes the primary upstream LSR as a candidate. Two streams of MPLS traffic are sent to the egress node from the two different upstream peers. The MPLS traffic from only one of the upstream neighbors is selected as the primary path to accept the traffic, and the other becomes the backup path. The traffic on the backup path is dropped. When the primary upstream path fails, the traffic from the backup path is then accepted. The multipoint LDP node selects the two upstream paths based on the interior gateway protocol (IGP) root node next hop.  Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
<b>Backup route</b>	For multipoint LDP with MoFRR, the route that is used if the primary route becomes unavailable.

## Sample Output

### show ldp traffic-statistics

```
user@host> show ldp traffic-statistics
```

FEC	Type	Packets	Bytes	Shared
10.35.3.0/30	Transit	0	0	Yes
	Ingress	0	0	No
10.35.10.1/32	Transit	0	0	Yes

	Ingress	0	0	No
10.255.245.214/32	Transit	0	0	No
	Ingress	11	752	No
192.168.37.36/30	Transit	0	0	Yes
	Ingress	0	0	No
FEC(root_addr:lsp_id)	Nexthop	Packets	Bytes	Shared
10.255.72.160:16777217	192.168.8.81	152056	14597376	No
	192.168.8.1	152056	14597376	No
	192.168.8.65	152056	14597376	No
NET FEC Statistics:				
FEC	Type	Packets	Bytes	Shared
10.255.107.230/32	Transit	30858	2022345	No
	Ingress	20	5120	No

#### show ldp traffic-statistics p2mp

```

user@host> show ldp traffic-statistics p2mp
FEC(root_addr:lsp_id) Nexthop      Packets      Bytes Shared
10.255.72.160:16777217 192.168.8.81  152056      14597376   No
                        192.168.8.1  152056      14597376   No
                        192.168.8.65  152056      14597376   No

```

#### show ldp traffic-statistics p2mp (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show ldp traffic-statistics p2mp
P2MP FEC Statistics:

FEC(root_addr:lsp_id/grp,src)  Nexthop      Packets      Bytes
Shared
11.99.0.73:239.10.0.1,11.98.0.10 11.99.0.117  243408      121217184
No
                        11.99.0.13    236286      117670428
No
11.99.0.73:239.10.0.2,11.98.0.10 11.99.0.117  248800      123902400
No
                        11.99.0.13    240759      119897982
No
11.99.0.73:239.10.0.1,11.98.0.20 11.99.0.117  250286      124642428
No
                        11.99.0.13    243741      121383018
No
11.99.0.73:239.10.0.2,11.98.0.20 11.99.0.117  252970      125979060
No
                        11.99.0.13    245218      122118564
No

```

#### show ldp traffic-statistics p2mp (Multipoint LDP with Multicast-Only Fast Reroute)

```

user@host> show ldp traffic-statistics p2mp

```

## P2MP FEC Statistics:

FEC(root_addr:lsp_id/grp,src)	Nexthop	Packets	Bytes
Shared			
1.1.1.1:232.1.1.1,192.168.219.11, Label: 301568	1.3.8.2	0	0
No	1.3.4.2	0	0
1.1.1.1:232.1.1.1,192.168.219.11, Label: 301584, Backup route	1.3.4.2	0	0
No	1.3.8.2	0	0
1.1.1.1:232.1.1.2,192.168.219.11, Label: 301600	1.3.8.2	0	0
No	1.3.4.2	0	0
1.1.1.1:232.1.1.2,192.168.219.11, Label: 301616, Backup route	1.3.4.2	0	0
No	1.3.8.2	0	0

## traceroute mpls ldp

---

**Syntax** `traceroute mpls <ldp> fec`  
`<destination>`  
`<detail>`  
`<exp>`  
`<fanout>`  
`<logical-system>`  
`<no-resolve>`  
`<paths>`  
`<retries>`  
`<routing-instance>`  
`<source>`  
`<ttl>`  
`<update>`  
`<wait>`

**Release Information** Command introduced in Junos OS Release 8.4.  
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Trace route to a remote host for an MPLS label-switched path signaled by the LDP. Use **traceroute mpls ldp** as a debugging tool to locate MPLS label-switched path forwarding issues in a network. (Currently supported for IPv4 packets only.)

**Options** *fec*—Specify the IP address and optional prefix of the forwarding equivalence class (FEC).  
*destination*—(Optional) Specify the destination address to use when sending probes.  
*detail*—(Optional) Display detailed output.  
*exp*—(Optional) Specify the class-of-service to use when sending probes. The range of values is 0 through 7. The default value is 7.  
*fanout*—(Optional) Specify the maximum number of nexthops to search per node. The range of values is 1 through 16. The default value is 16.  
*logical-system*—(Optional) Specify the name of the logical system for the traceroute attempt.  
*no-resolve*—(Optional) Specify not to resolve the hostname that corresponds to the IP address.  
*paths*—(Optional) Specify the number of paths to search. The range of values is 1 through 255. The default value is 16.  
*retries*—(Optional) Specify the number of times to resend probe. values. The range of values is 1 through 9. The default value is 3.  
*routing-instance routing-instance-name*—(Optional) Specify the name of the routing instance for the traceroute attempt.  
*source source-address*—(Optional) Specify the source address of the outgoing traceroute packets.

**ttl value**—(Optional) Specify the maximum time-to-live value to include in the traceroute request, in seconds. The range of values is **1** through **125** and the default value is **64**.

**update**—(Optional) Update database contents with traceroute results.

**wait seconds**—(Optional) Specify the number of seconds to wait before resending a probe. The range of values is **5** through **15** and the default value is **10** seconds.

**Required Privilege Level**

network

**List of Sample Output**

[traceroute mpls ldp on page 126](#)  
[traceroute mpls ldp detail on page 126](#)

**Output Fields**

[Table 14 on page 125](#) describes the output fields for the **traceroute mpls ldp fec** command and the **traceroute mpls ldp fec detail** commands. Output fields are listed in the approximate order in which they appear.

**Table 14: traceroute mpls ldp Output Fields**

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the <b>traceroute mpls ldp fec</b> command.	all levels
ttl	Time to live value of the labeled packet.	none specified
Label	Outgoing label used for forwarding the packet along the label-switched paths.	none specified
Protocol	Signaling protocol used. For this command, it is LDP.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is <b>null</b> .	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths).	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	<b>detail</b>
Parent	Address of the previous hop. Parent value for the first hop is <b>null</b> .	<b>detail</b>
Return Code	Return code for reporting the result of processing the echo request by the receiver.	<b>detail</b>
Response time	Time for the echo request to reach the receiver.	<b>detail</b>

Table 14: traceroute mpls ldp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is <b>none</b> .	<b>detail</b>
Label Stack	Label stack used to forward the packet.	<b>detail</b>

## Sample Output

### traceroute mpls ldp

```
user@router> traceroute mpls ldp 4.4.4.4
```

```
Probe options: ttl 64, retries 3, wait 10, paths 16, exp 7, fanout 16
ttl  Label Protocol Address Previous Hop Probe Status
  1   100016 LDP      24.24.24.1 (null) Success
  2   100000 LDP      20.20.20.2 24.24.24.1 Success
  3      3 LDP      22.22.22.4 20.20.20.2 Egress
```

```
Path 1 via fe-0/3/3.101 destination 127.0.0.64
```

### traceroute mpls ldp detail

```
user@router> traceroute mpls ldp 4.4.4.4 detail
```

```
Probe Options: ttl 64, retries 3, wait 10, paths 3, exp 7
Hop 24.24.24.1 Depth 1
  Parent (null)
  Return code: Label switched at stack-depth 1
  Response time 165.93 msec
  Multipath type: IP bitmask
  Address Range 1: 127.0.0.0 ~ 127.0.3.255
  Label Stack:
    Label 1 Value 100032 Protocol LDP

Hop 20.20.20.2 Depth 2
  Parent 24.24.24.1
  Return code: Upstream interface index unknown label-switched at stack-depth
1
  Response time 19.05 msec
  Multipath type: IP bitmask
  Address Range 1: 127.0.0.0 ~ 127.0.3.255
  Label Stack:
    Label 1 Value 100000 Protocol LDP

Hop 22.22.22.4 Depth 3
  Parent 20.20.20.2
  Return code: Egress-ok at stack-depth 1
  Response time 0.79 msec
  Multipath type: None
  Label Stack:
    Label 1 Value 3 Protocol LDP
```

## PART 2

# MPLS

- [Using MPLS on page 129](#)
- [Configuration Statements for MPLS on page 225](#)
- [Monitoring Commands for MPLS on page 315](#)





## CHAPTER 4

# Using MPLS

- [MPLS Overview For QFX Series and EX4600 Switches on page 130](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)
- [MPLS Limitations on QFX Series and EX4600 Switches on page 142](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 146](#)
- [Understanding MPLS Label Operations on page 150](#)
- [Understanding BGP on page 154](#)
- [IPv6 Layer 3 VPNs on page 156](#)
- [Ethernet Pseudowire Overview on page 157](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 158](#)
- [Understanding Ethernet-over-MPLS \(L2 Circuit\) on page 161](#)
- [Understanding Using MPLS-Based Layer 3 VPNs on Switches on page 162](#)
- [Carrier-of-Carriers VPNs on page 164](#)
- [Interprovider and Carrier-of-Carriers VPNs on page 165](#)
- [Chained Composite Next Hops for Transit Devices for VPNs on page 166](#)
- [Fast Reroute Overview on page 168](#)
- [Graceful Restart and MPLS-Related Protocols on page 170](#)
- [Types of LSPs on page 171](#)
- [Configuring Automatic Bandwidth Allocation for LSPs on page 172](#)
- [Configuring CoS Bits for an MPLS Network on page 179](#)
- [Configuring Ethernet over MPLS \(L2 Circuit\) on page 180](#)
- [Configuring a Global MPLS EXP Classifier on page 184](#)
- [Configuring MPLS Firewall Filters and Policers on page 184](#)
- [Configuring MPLS to Gather Statistics on page 187](#)
- [Configuring MPLS on Provider Edge Switches on page 188](#)
- [Configuring MPLS on Provider Switches on page 192](#)
- [Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs on page 193](#)
- [Configuring Static Label Switched Paths for MPLS on page 196](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 199](#)

- [Example: Configuring MPLS-Based Layer 3 VPNs on page 200](#)
- [Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks on page 209](#)
- [Verifying That MPLS Is Working Correctly on page 218](#)
- [MPLS Configuration Guidelines on page 220](#)
- [Supported MPLS Scaling Values on page 221](#)
- [MPLS Stitching For Virtual Machine Connection on page 222](#)

---

## MPLS Overview For QFX Series and EX4600 Switches

---

Multiprotocol Label Switching (MPLS) is a protocol that uses labels to route packets instead of using IP addresses. In a traditional network, each switch performs an IP routing lookup, determines a next-hop based on its routing table, and then forwards a packet to that next-hop. With MPLS, only the first device does a routing lookup, and, instead of finding the next-hop, finds the ultimate destination along with a path to that destination. The path of an MPLS packet is called a label-switched path (LSP).

MPLS applies one or more labels to a packet so it can follow the LSP to the destination. Each switch pops off its label and sends the packet to the next switch label in the sequence.

The Junos OS includes everything you need to configure MPLS. You do not need to install any additional programs or protocols. MPLS is supported on switches with a subset of the commands supported on routers. The Junos MPLS-configured switches can interact with each other and with Junos MPLS-configured routers.

MPLS has the following advantages over conventional packet forwarding:

- Packets arriving on different ports can be assigned different labels.
- A packet arriving at a particular provider edge (PE) switch can be assigned a label that is different from that of the same packet entering the network at a different PE switch. As a result, forwarding decisions that depend on the ingress PE switch can be easily made.
- Sometimes it is desirable to force a packet to follow a particular route that is explicitly chosen at or before the time the packet enters the network, rather than letting it follow the route chosen by the normal dynamic routing algorithm as the packet travels through the network. In MPLS, a label can be used to represent the route so that the packet need not carry the identity of the explicit route.

This topic describes:

- [Why Use MPLS? on page 131](#)
- [Why Not Use MPLS? on page 131](#)
- [How Do I Configure MPLS? on page 131](#)
- [What Does the MPLS Protocol Do? on page 132](#)
- [How Does MPLS Interface to Other Protocols? on page 133](#)
- [If I Have Used Cisco MPLS, What Do I Need to Know? on page 133](#)

## Why Use MPLS?

MPLS reduces the use of the forwarding table by using labels instead of the forwarding table. The size of forwarding tables on a switch are limited by silicon and using exact matching for forwarding to destination devices is cheaper than buying more sophisticated hardware. In addition, MPLS allows you to control where and how traffic is routed on your network – this is called traffic engineering.

Some reasons to use MPLS instead of another switching solution are:

- MPLS can connect different technologies that would not otherwise be compatible---service providers have this compatibility issue when connecting clients with different autonomous systems in their networks. In addition, MPLS has a feature called Fast Reroute that provides alternate backups for paths – this prevents network degradation in case of a switch failure.
- Other IP-based encapsulations such as Generic Route Encapsulation (GRE) or Virtual Extensible Local Area Networks (VXLAN) support only two levels of hierarchy, one for the transport tunnel and one piece of metadata. Using virtual servers means that you need multiple hierarchy levels. For example, one label is needed for top-of-rack (ToR), one label for the egress port that identifies the server, and one for the virtual server.

## Why Not Use MPLS?

There are no protocols to auto-discover MPLS enabled nodes. MPLS protocol just exchanges label values for an LSP. They do not create the LSPs.

You must build the MPLS mesh, switch by switch. We recommend using scripts for this repetitive process.

MPLS hides suboptimal topologies from BGP where multiple exits may exist for the same route.

Large LSPs are limited by the circuits they traverse. You can work around this by creating multiple, parallel LSPs.

## How Do I Configure MPLS?

There are three types of switches you must set up for MPLS:

- Label Edge Router/Switch (LER) or ingress node to the MPLS network. This switch encapsulates the packets.
- Label Switching Routers/Switches (LSR). One or more switches that transfer MPLS packets in the MPLS network.
- Egress router/switch is the final MPLS device that removes the last label before packets leave the MPLS network.

Service providers (SP) use the term provider router (P) for a backbone router/switch doing label switching only. The customer-facing router at the SP is called a provider edge router (PE). Each customer needs a customer edge router (CE) to communicate with

the PE. Customer facing routers typically can terminate IP addresses, L3VPNs, L2VPNs/pseudowires, and VPLS before packets are transferred to the CE.

### **Configure the MPLS LER (Ingress) Switch and the Egress Switch**

---

To configure MPLS, you must first create one or more named paths on the ingress and egress routers. For each path, you can specify some or all transit routers in the path, or you can leave it empty. See *Configuring the Ingress Router for MPLS-Signaled LSPs* and *Configuring the Intermediate and Egress Routers for MPLS-Signaled LSPs*, *Configuring the Ingress and Egress Router Addresses for LSPs*, and *Configuring the Connection Between Ingress and Egress Routers*.

### **Configure LSRs for MPLS**

---

Configure one or more MPLS LSRs by following these steps:

1. Configure interfaces on each switch to transmit and receive MPLS packets using the usual interface command with MPLS appended. For example:

```
[edit interfaces ge-0/0/0 unit 0] family mpls;
```

2. Add those same interfaces under [edit protocols mpls]. For example:

```
[edit protocols mpls]
  interface ge-0/0/0;
```

3. Configure the interfaces on each switch to handle MPLS labels with a protocol. For example, for LDP:

```
[edit protocols ldp]
  Interface ge-0/0/0.0;
```

To watch a demo of these configurations, see  
<https://www.youtube.com/watch?v=xegWBCUJ4tE>.

## **What Does the MPLS Protocol Do?**

Multiprotocol Label Switching (MPLS) is an Internet Engineering Task Force (IETF)-specified framework that provides for the designation, routing, forwarding and switching of traffic flows through the network. In addition, MPLS:

- Specifies mechanisms to manage traffic flows of various granularities, such as flows between different hardware, machines, or even flows between different applications.
- Remains independent of the layer-2 and layer-3 protocols.
- Provides a means to map IP addresses to simple, fixed-length labels used by different packet-forwarding and packet-switching technologies.
- Interfaces to existing routing protocols, such as Resource ReSerVation Protocol (RSVP) and Open Shortest PathFirst (OSPF).
- Supports IP, ATM, and Frame Relay layer-2 protocols.
- Uses these additional technologies:

- FRR: MPLS Fast Reroute improves convergence during a failure by mapping out alternate LSPs in advance.
- Link Protection/ Next-hop backup: A bypass LSP is created for every possible link failure.
- Node Protection/ Next-hop backup: A bypass LSP is created for every possible switch (node) failure.
- VPLS: Creates Ethernet multipoint switching service over MPLS and emulates functions of an L2 switch.
- L3VPN: IP-based VPN customers get individual virtual routing domains.

## How Does MPLS Interface to Other Protocols?

Some of the protocols that work with MPLS are:

- RSVP-TE: Resource Reservation Protocol - Traffic Engineering reserves bandwidth for LSPs.
- LDP: Label Distribution Protocol is the defacto protocol used for distribution of MPLS packets and is usually configured to tunnel inside RSVP-TE.
- IGP: Interior Gateway Protocol is a routing protocol. Edge routers (PE-routers) run BGP between themselves to exchange external (customer) prefixes. Edge and core (P) routers run IGP (usually OSPF or IS-IS) to find optimum path toward BGP next hops. P- and PE-routers use LDP to exchange labels for known IP prefixes (including BGP next hops). LDP indirectly builds end-to-end LSPs across the network core.
- BGP: Border Gateway Protocol (BGP) is allows policy-based routing to take place, using TCP as its transport protocol on port 179 to establish connections. The Junos OS routing protocol software includes BGP version 4. You do not configure BGP---configuring interfaces with MPLS and LDP/RSVP establishes the labels and the ability to transmit packets. BGP automatically determines the routes packets take.
- OSPF and ISIS: These protocols are used for routing between the MPLS PE and CE. Open Shortest Path First (OSPF) is perhaps the most widely used interior gateway protocol (IGP) in large enterprise networks. IS-IS, another link-state dynamic routing protocol, is more common in large service provider networks. Assuming you're running L3VPN to your customers, on the SP edge between the PE and the CE you can run any protocol that your platform supports as a VRF aware instance.

## If I Have Used Cisco MPLS, What Do I Need to Know?

Cisco Networks and Juniper Networks use different MPLS terminology.

What Cisco Calls:	Juniper Calls:
affinities	admin-groups
autoroute announce	TE shortcuts
forwarding adjacency	LSP-advertise

What Cisco Calls:	Juniper Calls:
tunnel	LSP
make-before-break	adaptive
application-window	adjust-interval
shared risk link groups	fate-sharing

**Related Documentation**

- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 146](#)
- [Understanding MPLS Label Operations on page 150](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 158](#)
- [Junos OS MPLS Applications Library for Routing Devices](#)

## MPLS Feature Support on QFX Series and EX4600 Switches

Multiprotocol Label Switching (MPLS) is a set of procedures for augmenting network layer packets with label stacks, thereby turning them into labeled packets. Service providers frequently use MPLS. Simply put, where traditional networks' routers each perform an IP lookup to determine the next hop, an MPLS network's first device does a routing lookup for the final destination instead of the next hop. A label is then applied to the packet—this is called packet switching. The final destination device removes the label.

A number of Juniper Networks switches are capable of running a subset of MPLS and can therefore communicate, not only with each other, but with Juniper Networks routers running MPLS. This topic describes the major MPLS features that are supported on QFX Series switches and on one EX Series switch, the EX4600. Be sure to check for any exceptions to this support in "[MPLS Limitations on QFX Series and EX4600 Switches](#)" on page 142.



**NOTE:** EX4600 switches use the same chipset as QFX5100 switches—this is why that EX Series switch is discussed here along with QFX Series switches. Other EX Series switches also support MPLS but with a different feature set.

This topic describes:

- [MPLS Commands Supported by QFX Series and EX4600 Switches on page 135](#)
- [MPLS Features Supported by QFX Series and EX4600 Switches on page 135](#)

## MPLS Commands Supported by QFX Series and EX4600 Switches

QFX Series and EX4600 switches support a subset of MPLS features. The command-line interface (CLI) for switches displays all MPLS related configuration statements, even those that are not supported. However, configuring those unsupported statements on a switch has no effect on the operation of the switch.

## MPLS Features Supported by QFX Series and EX4600 Switches

This section lists the major MPLS features supported on QFX Series and EX4600 switches and the Junos OS release in which they were introduced. [Table 15 on page 135](#) lists the features for the QFX10000 Series switches. [Table 16 on page 137](#) lists the features for the QFX3500, QFX5100, and QFX5200 switches. [Table 17 on page 140](#) lists the features for the EX4600 switches.

**Table 15: QFX10000 Switch MPLS Features with Junos OS Release Support**

Feature	QFX10002	QFX10008	QFX10016
QFX standalone switch as an MPLS provider edge (PE) switch or provider switch	15.1X53-D10	15.1X53-D30	15.1X53-D60
Label edge router (LER)	15.1X53-D10	15.1X53-D30	15.1X53-D60
Label switch router (LSR)	15.1X53-D10	15.1X53-D30	15.1X53-D60
Route reflector for BGP labeled routes	15.1X53-D10	15.1X53-D30	15.1X53-D60
Auto-bandwidth	15.1X53-D60	15.1X53-D60	15.1X53-D60
Border Gateway Protocol (BGP) labeled unicast	15.1X53-D10	15.1X53-D30	15.1X53-D60
Carrier-over-carrier and inter-provider BGP L3 VPN	not supported	not supported	not supported
Class of Service (CoS or QoS) for MPLS traffic	15.1X53-D10	15.1X53-D30	15.1X53-D60
Ethernet-over-MPLS (L2 circuit)	15.1X53-D60	15.1X53-D60	15.1X53-D60
Fast Reroute (FRR), one-to-one local protection and many-to-one local protection	15.1X53-D10	15.1X53-D30	15.1X53-D60
FRR using detours and secondary LSP	15.1X53-D10	15.1X53-D30	15.1X53-D60
Firewall filters	15.1X53-D30	15.1X53-D30	15.1X53-D60

Table 15: QFX10000 Switch MPLS Features with Junos OS Release Support (*continued*)

Feature	QFX10002	QFX10008	QFX10016
Graceful restart for Open Shortest Path First (OSPF) for Resource Reservation Protocol (RSVP)	15.1X53-D10	15.1X53-D30	15.1X53-D60
Intermediate System to Intermediate System (ISIS) routing protocol as an interior gateway protocol (IGP) for MPLS. IS-IS interior gateway protocol traffic engineering (TE)	15.1X53-D10	15.1X53-D30	15.1X53-D60
IP-over-MPLS label-switched paths (LSPs) both static and dynamic links	15.1X53-D10	15.1X53-D30	15.1X53-D60
IPv6 tunneling for over an MPLS-based IPv4 network (6PE)  Layer 3 VPN 6PE	15.1X53-D10	15.1X53-D30	15.1X53-D60
Label Distribution Protocol (LDP) based signaling over RSVP	15.1X53-D10	15.1X53-D30	15.1X53-D60
Layer 3 VPNs IPv4	15.1X53-D10	15.1X53-D30	15.1X53-D60
MPLS over integrated bridging and routing (IRB) interfaces	15.1X53-D10	15.1X53-D30	15.1X53-D60
MTU signaling in RSVP	15.1X53-D10	15.1X53-D30	15.1X53-D60
Object access method (OAM) including MPLS ping, traceroute and BFD	15.1X53-D10	15.1X53-D30	15.1X53-D60
Open Shortest Path First traffic engineering (OSPF TE)	15.1X53-D10	15.1X53-D30	15.1X53-D60
OSPFv2 as an interior gateway protocol	15.1X53-D10	15.1X53-D30	15.1X53-D60
Pseudowire-over-aggregated Ethernet interfaces (core-facing interface)	15.1X53-D60 (supported only on NNI interfaces)	15.1X53-D60 (supported only on NNI interfaces)	15.1X53-D60 (supported only on NNI interfaces)
Resource Reservation Protocol (RSVP)	15.1X53-D10	15.1X53-D30	15.1X53-D60
RSVP bandwidth	15.1X53-D10	15.1X53-D30	15.1X53-D60
RSVP fast reroute including link-protection, node-link-protection, FRR using detours, and secondary LSP	15.1X53-D10	15.1X53-D30	15.1X53-D60
RSVP Traffic engineering (used to establish LSPs) with IS-IS and OSPF extensions	15.1X53-D10	15.1X53-D30	15.1X53-D60



Table 15: QFX10000 Switch MPLS Features with Junos OS Release Support (*continued*)

Feature	QFX10002	QFX10008	QFX10016
SNMP MIB support	15.1X53-D10	15.1X54-D30	15.1X53-D60
Static and dynamic LSPs	15.1X53-D10	15.1X53-D30	15.1X53-D60
Traffic Engineering (TE)	15.1X53-D10	15.1X53-D30	15.1X53-D60
TE auto-bandwidth and RSVP bandwidth	15.1X53-D10	15.1X53-D30	15.1X53-D60
Dynamic bandwidth management using ingress LSP splitting and merging			
Virtual routing and forwarding (VRF) label support	15.1X53-D10	15.1X53-D30	15.1X53-D60

Table 16: QFX3500, QFX5100, and QFX5200 MPLS Features with Junos OS Release Support

Feature	QFX3500	QFX5100	QFX5200
QFX standalone switch as an MPLS provider edge (PE) switch or provider switch	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D30
Label edge router (LER)	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D30
Label switch router (LSR)	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D30
Route reflector for BGP labeled routes	15.1X53-D10	15.1X53-D30	15.1X53-D30
Auto-bandwidth	not supported	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D30
Border Gateway Protocol (BGP) labeled unicast	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D30
Carrier-over-carrier and inter-provider BGP L3 VPN	14.1X53-D15	14.1X53-D15 VC/VCF (14.1X53-D30)	15.1X53-D30

Table 16: QFX3500, QFX5100, and QFX5200 MPLS Features with Junos OS Release Support (*continued*)

Feature	QFX3500	QFX5100	QFX5200
Class of Service (CoS or QoS) for MPLS traffic	12.3X50-D10	13.2X51-D15  VC/VCF (14.1X53-D30)	15.1X53-D30
ECMP at LSR routers: <ul style="list-style-type: none"> <li>• SWAP</li> <li>• PHP</li> <li>• L3VPN</li> <li>• L2 Circuit</li> </ul>	not supported	14.1X53-D35 (Supported only on label stack. Not supported on flow label, entropy label, or ECMP label)	15.1X53-D30
Entropy labels	not supported	not supported	not supported
Ethernet-over-MPLS (L2 circuit)	14.1X53-D10	14.1X53-D10  VC/VCF (14.1X53-D30)	15.1X53-D30
Fast Reroute (FRR), one-to-one local protection and many-to-one local protection	14.1X53-D10	14.1X53-D10  VC/VCF (not supported)	15.1X53-D30
FRR using detours and secondary LSP	not supported	not supported	not supported
Firewall filters	12.3X50-D10	13.2X51-D15  VC/VCF (14.1X53-D30)	15.1X53-D30
Flow-Aware Transport of Pseudowires (FAT) Flow Labels	not supported	not supported  VC/VCF (not supported)	not supported
Graceful restart for Open Shortest Path First (OSPF) for Resource Reservation Protocol (RSVP)	12.2X50-D10	13.2X51-D15  VC/VCF (14.1X53-D30)	15.1X53-D30
Intermediate System to Intermediate System (ISIS) routing protocol as an interior gateway protocol (IGP) for MPLS. IS-IS interior gateway protocol traffic engineering (TE)	12.2X50-D10	13.2X51-D15  VC/VCF (14.1X53-D30)	15.1X53-D30
IP-over-MPLS label-switched paths (LSPs) both static and dynamic links	12.2X50-D10	13.2X51-D15  VC/VCF (14.1X53-D30)	15.1X53-D30

**Table 16: QFX3500, QFX5100, and QFX5200 MPLS Features with Junos OS Release Support (*continued*)**

Feature	QFX3500	QFX5100	QFX5200
IPv6 tunneling for over an MPLS-based IPv4 network (6PE)	12.3X50-D10	13.2X51-D15	15.1X53-D30
Layer 3 VPN 6PE		VC/VCF (14.1X53-D30)	
Label Distribution Protocol (LDP) based signaling over RSVP	12.2X50-D10	13.2X51-D15  VC/VCF (14.1X53-D30)	15.1X53-D30
Layer 3 VPNs IPv4	12.3X50-D10	13.2X51-D15  VC/VCF (14.1X53-D30)	15.1X53-D30
Loop-Free Alternate (LFA)	not supported	13.2X51-D15  VC/VCF (14.1X53-D30)	not supported
MPLS over integrated bridging and routing (IRB) interfaces	not supported	not supported	not supported
MTU signaling in RSVP	12.3X50-D10	13.2X51-D15  VC/VCF (14.1X53-D30)	15.1X53-D30
Object access method (OAM) including MPLS ping, traceroute and BFD	12.3X50-D10	13.2X51-D15  VC/VCF (14.1X53-D30)	15.1X53-D30
Open Shortest Path First traffic engineering (OSPF TE)	12.3X50-D10	13.2X51-D15  VC/VCF (not supported)	15.1X53-D30
OSPFv2 as an interior gateway protocol	12.2X50-D10	13.2X51-D15  VC/VCF (14.1X53-D30)	15.1X53-D30
Pseudowire-over-aggregated Ethernet interfaces (core-facing interface)	14.1X53-D10	14.1X53-D15  VC/VCF (14.1X53-D30)	15.1X53-D30
Resource Reservation Protocol (RSVP), bandwidth and auto-bandwidth	12.2X50-D10	13.2X51-D15  VC/VCF (14.1X53-D30)	15.1X53-D30

**Table 16: QFX3500, QFX5100, and QFX5200 MPLS Features with Junos OS Release Support (*continued*)**

Feature	QFX3500	QFX5100	QFX5200
RSVP fast reroute including link-protection, node-link-protection, FRR using detours, and secondary LSP	14.1X53-D15	14.1X53-D15 VC/VCF (not supported)	15.1X53-D30
RSVP Traffic engineering (used to establish LSPs) with IS-IS and OSPF extensions	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D30
SNMP MIB support	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D30
Static and dynamic LSPs	12.2X50-D10	13.2X51-D10 VC/VCF (14.1X53-D30)	15.1X53-D30
Traffic Engineering (TE)	13.1X51-D10	13.1X51-D10	15.1X53-D30
TE auto-bandwidth and RSVP bandwidth	13.1X51-D10	13.1X51-D10	15.1X53-D30
Virtual routing and forwarding (VRF) label support	12.2X50-D10	13.2X51-D15 VC/VCF (14.1X53-D30)	15.1X53-D30

**Table 17: EX4600 Switch MPLS Features with Junos OS Release Support**

Feature	EX4600
QFX standalone switch as an MPLS provider edge (PE) switch or provider switch	14.1X53-D15 but not 15.1X53-D10
Label edge router (LER)	14.1X53-D15 but not 15.1X53-D10
Label switch router (LSR)	14.1X53-D15 but not 15.1X53-D10
Route reflector for BGP labeled routes	14.1X53-D15 but not 15.1X53-D10
Border Gateway Protocol (BGP) labeled unicast	14.1X53-D15 but not 15.1X53-D10
Carrier-over-carrier and inter-provider BGP L3 VPN	not supported
Class of Service (CoS or QoS) for MPLS traffic	14.1X53-D15 but not 15.1X53-D10

Table 17: EX4600 Switch MPLS Features with Junos OS Release Support (*continued*)

Feature	EX4600
Ethernet-over-MPLS (L2 Circuit)	14.1X53-D15 but not 15.1X53-D10
Fast Reroute (FRR), one-to-one local protection and many-to-one local protection	not supported
FRR using detours and secondary LSP	not supported
Firewall filters	14.1X53-D15 but not 15.1X53-D10
Graceful restart for Open Shortest Path First (OSPF) for Resource Reservation Protocol (RSVP)	13.2X51-D25
IP-over-MPLS label-switched paths (LSPs) both static and dynamic links	14.1X53-D15 but not 15.1X53-D10
IPv6 tunneling for over an MPLS-based IPv4 network (6PE)	14.1X53-D15 but not 15.1X53-D10
Layer 3 VPN 6PE	
Intermediate System to Intermediate System (ISIS) routing protocol as an interior gateway protocol (IGP) for MPLS. IS-IS interior gateway protocol traffic engineering (TE)	14.1X53-D15 but not 15.1X53-D10
Label Distribution Protocol (LDP) based signaling over RSVP	14.1X53-D15 but not 15.1X53-D10
Layer 3 VPNs IPv4	14.1X53-D15 but not 15.1X53-D10
MPLS over integrated bridging and routing (IRB) interfaces	not supported
MTU signaling in RSVP	14.1X53-D15 but not 15.1X53-D10
Object access method (OAM) including MPLS ping, traceroute and BFD	14.1X53-D15 but not 15.1X53-D10
Open Shortest Path First traffic engineering (OSPF TE)	14.1X53-D15 but not 15.1X53-D10
OSPFv2 as an interior gateway protocol	13.2X51-D25
Pseudowire-over-aggregated Ethernet interfaces (core-facing interface)	14.1X53-D15 but not 15.1X53-D10

Table 17: EX4600 Switch MPLS Features with Junos OS Release Support (*continued*)

Feature	EX4600
Resource Reservation Protocol (RSVP), bandwidth and auto-bandwidth	14.1X53-D15 but not 15.1X53-D10
RSVP fast reroute including link-protection, node-link-protection, FRR using detours, and secondary LSP	not supported
RSVP Traffic engineering (used to establish LSPs) with IS-IS and OSPF extensions	14.1X53-D15 but not 15.1X53-D10
SNMP MIB support	14.1X53-D15 but not 15.1X53-D10
Static and dynamic LSPs	14.1X53-D15 but not 15.1X53-D10
Traffic Engineering (TE)	14.1X53-D15 but not 15.1X53-D10
TE auto-bandwidth and RSVP bandwidth	14.1X53-D15 but not 15.1X53-D10
Virtual routing and forwarding (VRF) label support	14.1X53-D15 but not 15.1X53-D10

- Related Documentation**
- [MPLS Limitations on QFX Series and EX4600 Switches on page 142](#)
  - [MPLS Configuration Guidelines on page 220](#)

## MPLS Limitations on QFX Series and EX4600 Switches

MPLS is fully implemented on routers, while switches support a subset of the MPLS features. The limitations of each switch are listed in a separate section here, even though many of the limitations are duplicates that apply to more than one switch.

- [MPLS Limitations on QFX3500 Switches on page 143](#)
- [MPLS Limitations on QFX5100 and EX4600 Switches on page 143](#)
- [MPLS Limitations on QFX5100 Virtual Chassis and Virtual Chassis Fabric on page 145](#)
- [MPLS Limitations on QFX10000 Switches on page 145](#)

## MPLS Limitations on QFX3500 Switches

- If you configure the BGP labeled unicast address family (using the **labeled-unicast** statement at the **[edit protocols bgp family inet]** hierarchy level) on a QFX switch or on an EX4600 switch deployed as a route reflector for BGP labeled routes, path selection will occur at the route reflector, and a single best path will be advertised. This will result in loss of BGP multipath information.
- Fast reroute is supported, however the **include-all** and **include-any** options for fast reroute are not supported. For more information, see [“Fast Reroute Overview” on page 168](#).
- MPLS-based circuit cross-connects (CCC) are not supported—only circuit-based pseudowires are supported.
- MTU signaling in RSVP and discovery is supported in the Control Plane. However, this cannot be enforced in data plane.
- With L2 circuit-based pseudowires, if multiple equal-cost RSVP LSP's are available to reach a Layer 2 Circuit neighbor, one LSP is randomly used for forwarding. Use this feature to specify LSPs for specific L2 circuit traffic to load-share the traffic in the MPLS core.
- Configuring an MPLS firewall filter on a switch that is deployed as an egress provider edge (PE) switch has no effect.
- Configuring the **revert-timer** statement at the **[edit protocols mpls]** hierarchy level has no effect.

## MPLS Limitations on QFX5100 and EX4600 Switches

- On a QFX5100 switch, you can observe traffic drop after changing your configuration to enable VLAN tagged for MPLS packets. As a result of packet capture, a QFX5100 switch can swap the wrong VLAN ID for MPLS packets.
- Even though both QFX5100 and EX4600 switches use the same chipset, MPLS support differs. EX4600 switches support only basic MPLS functionality while QFX5100 switches support some of the more advanced features. See [“MPLS Feature Support on QFX Series and EX4600 Switches” on page 134](#) for details.
- If you configure the BGP labeled unicast address family (using the **labeled-unicast** statement at the **[edit protocols bgp family inet]** hierarchy level) on a QFX switch or on an EX4600 switch deployed as a route reflector for BGP labeled routes, path selection will occur at the route reflector, and a single best path will be advertised. This will result in loss of BGP multipath information.
- Fast reroute is supported, however the **include-all** and **include-any** options for fast reroute are not supported. For more information, see [“Fast Reroute Overview” on page 168](#).
- MPLS-based circuit cross-connects (CCC) are not supported—only circuit-based pseudowires are supported.

- MTU signaling in RSVP and discovery is supported in the Control Plane. However, this cannot be enforced in data plane.
- With L2 circuit-based pseudowires, if multiple equal-cost RSVP LSP's are available to reach a Layer 2 Circuit neighbor, one LSP is randomly used for forwarding. Use this feature to specify LSPs for specific L2 circuit traffic to load-share the traffic in the MPLS core.
- Configuring an MPLS firewall filter on a switch that is deployed as an egress provider edge (PE) switch has no effect.
- Configuring the **revert-timer** statement at the **[edit protocols mpls]** hierarchy level has no effect.
- Equal-cost multi-path routing (ECMP) for MPLS is not supported on the QFX5100 switch.
- These are hardware limitations for EX4600 and QFX5100 switches:
  - Push of a maximum of 3 labels is supported in the MPLS edge switch if label swap is not done.
  - Push of a maximum of 2 labels is supported in the MPLS edge switch if label swap is done.
  - Pop at line rate is supported for a maximum of 2 labels.
  - Global label space is supported but interface-specific label space is not supported.
  - ECMP based on incoming labels at LSR is not supported.
  - QFX switches with Broadcom chips do not support separate next hops for the same label with different S bits (S-0 and S-1). This includes the QFX3500, QFX3600, QFX5100, and QFX5200 switches.
  - On the QFX5100, the MPLS MTU command can cause unexpected behavior—this is due to SDK chipset limitations on this platform.
- These LDP features are not supported on the QFX5100 switches:
  - LDP multipoint
  - LDP link protection
  - LDP bidirectional forwarding detection (BFD)
  - LDP operation administration and management (OAM)
  - LDP multicast-only fast reroute (MoFRR)
  - LDP equal-cost multipath (ECMP)
- On the QFX5100, IRB/L3-sub interfaces are not supported on the NNI port in an L2 circuit configuration. Because of this, a QFX5100 is unable to ping the neighbor IP of either a direct or local interface.



## MPLS Limitations on QFX5100 Virtual Chassis and Virtual Chassis Fabric

The following MPLS features are not supported by the QFX5100 VC and QFX5100 VCF switches:

- Next-hop LSP Details in section 2.2.6.1
- BFD including BFD triggered FRR
- L2VPN based on BGP (VPWS / draft Kompella)
- VPLS
- Extended-vlan-ccc
- Pseudo-wire protection using Ethernet OAM
- Local switching of pseudo-wire
- Pseudowire fault detection based on VCCV
- QFX switches with Broadcom chips do not support separate next hops for the same label with different S bits (S-0 and S-1). This includes QFX3500, QFX3600, QFX5100, and QFX5200 switches.

## MPLS Limitations on QFX10000 Switches

- Carrier-of-carriers and inter-provider VPNs are not needed on QFX10000 Series switches.
- Configuring an MPLS firewall filter on a switch that is deployed as an egress provider edge (PE) switch has no effect.
- Configuring the **revert-timer** statement at the **[edit protocols mpls]** hierarchy level has no effect.
- These LDP features are not supported on the QFX10000 switches:
  - LDP multipoint
  - LDP link protection
  - LDP bidirectional forwarding detection (BFD)
  - LDP operation administration and management (OAM)
  - LDP multicast-only fast reroute (MoFRR)
- Pseudowire-over-aggregated Ethernet interfaces on UNI are not supported.

### Related Documentation

- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)

## Understanding MPLS Components for QFX Series and EX4600 Switches

---

MPLS devices include a number of components. While some components are required for all MPLS applications, others might not be, depending on the specific application.

This topic includes:

- [Provider Edge Switches on page 146](#)
- [Provider Switch on page 147](#)
- [Components Required for All Switches in the MPLS Network on page 148](#)

### Provider Edge Switches

To implement MPLS on a network, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch. In addition, you must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets.

The ingress PE switch (the entry point to the MPLS tunnel) receives a packet, analyzes it, and pushes an MPLS label onto it. This label places the packet in a forwarding equivalence class (FEC) and determines its handling and destination through the MPLS tunnel. The egress PE switch (the exit point from the MPLS tunnel) pops the MPLS label off the outgoing packet.

Within an MPLS tunnel, the network traffic is bidirectional. Therefore, each PE switch can be configured to be both an ingress switch and an egress switch, depending on the direction of the traffic.

The following MPLS components are configured on the PE switches but not on the provider switches:

- [MPLS Protocol and Label-Switched Paths on page 146](#)
- [IP Over MPLS for Customer Edge Interfaces on page 146](#)
- [BGP Layer 3 VPN Configuration on page 147](#)
- [Routing Instances for Layer 3 VPN on page 147](#)
- [Routing Instances for Layer 2 VPN and Layer 3 VPN on page 147](#)
- [Ethernet Encapsulation for Layer 2 VPN on page 147](#)

### MPLS Protocol and Label-Switched Paths

---

Each PE switch must be configured to support the MPLS protocol. You must also configure label-switched paths (LSPs) at the **[edit protocols mpls]** hierarchy level.

### IP Over MPLS for Customer Edge Interfaces

---

You can configure the customer edge interfaces of the PE switches for IP over MPLS using a Layer 3 interface and a static route from the ingress PE switch to the egress PE switch. See “[Configuring MPLS on Provider Edge Switches](#)” on page 188.

### BGP Layer 3 VPN Configuration

---

If you are implementing a Layer 3 virtual private network (VPN), you must configure the BGP routing protocol on the PE switches.

### Routing Instances for Layer 3 VPN

---

If you are implementing a Layer 3 VPN, you must configure a routing instance. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables.

QFX Series and EX4600 devices support VPN routing and forwarding (VRF) routing instances for Layer 3 VPNs.

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table will be **my-instance.inet.0**. All routes for **my-instance** are installed in **my-instance.inet.0**.

### Routing Instances for Layer 2 VPN and Layer 3 VPN

---

If you are implementing a Layer 2 VPN or a Layer 3 VPN, you must configure a routing instance. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables.

QFX Series devices support the following types of routing instances:

- Layer 2 VPN—To support a Layer 2 VPN
- VPN routing and forwarding (VRF)—To support a Layer 3 VPN

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table will be **my-instance.inet.0**. All routes for **my-instance** are installed in **my-instance.inet.0**.

### Ethernet Encapsulation for Layer 2 VPN

---

If you are implementing a Layer 2 VPN, you must also configure the physical layer encapsulation type on the customer edge interface and within the routing instance.

## Provider Switch

You must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets. You can add provider switches without changing the configuration of the PE switches.

A provider switch does not analyze packets. It refers to an MPLS label forwarding table and swaps one label for another. The new label determines the next hop along the MPLS tunnel. A provider switch cannot perform push or pop operations.

## Components Required for All Switches in the MPLS Network

The following MPLS components are configured on both the PE switches and the provider switches:

- [Interior Gateway Protocol on page 148](#)
- [Traffic Engineering on page 148](#)
- [MPLS Protocol on page 148](#)
- [RSVP on page 148](#)
- [Family mpls on page 149](#)

---

### Interior Gateway Protocol

MPLS works in coordination with OSPF as the interior gateway protocol (IGP). Therefore, you must configure OSPF as the IGP on the loopback interface and CE-facing interfaces of both the PE switches and the provider switches.

The CE-facing interfaces can be either Gigabit Ethernet or 10-Gigabit Ethernet interfaces, and they can be configured as either individual interfaces or as aggregated Ethernet interfaces.



**NOTE:** The CE-facing interfaces cannot be configured with VLAN tagging or a VLAN ID. When you configure them to belong to family mpls, they are removed from the default VLAN if they were members of that VLAN. They operate as an exclusive tunnel for MPLS traffic.

---

---

### Traffic Engineering

Traffic engineering maps traffic flows onto an existing physical topology and provides the ability to move traffic flow away from the shortest path selected by the IGP and to a potentially less congested physical path across a network.

Traffic engineering enables the selection of specific end-to-end paths to send given types of traffic through your network. You must configure OSPF traffic engineering on the PE switches and the provider switches.

---

### MPLS Protocol

You must enable the MPLS protocol on all switches that participate in the MPLS network and apply it to the core interfaces of both the PE and provider switches. You do not need to apply it to the loopback interface because the MPLS protocol uses the framework established by the RSVP signaling protocol to create LSPs. On the PE switches, the configuration of the MPLS protocol must also include the definition of an LSP.

---

### RSVP

RSVP is a signaling protocol that allocates and distributes labels throughout an MPLS network. RSVP sets up unidirectional paths between the ingress PE switch and the egress

PE switch. RSVP makes the LSPs dynamic; it can detect topology changes and outages and establish new LSPs to allow traffic to move around a failure.

You must enable RSVP and apply it to the loopback interface and the core interface of both the PE and provider switches. The path message contains the configured information about the resources required for the LSP to be established.

When the egress PE switch receives the path message, it sends a reservation message back to the ingress PE switch. This reservation message is passed along from switch to switch along the same path as the original path message. Once the ingress PE switch receives this reservation message, an RSVP path is established.

The established LSP stays active as long as the RSVP session remains active. RSVP continues activity through the transmissions and responses to RSVP path and reservation messages. If the messages stop for three minutes, the RSVP session terminates and the LSP is lost.

RSVP runs as a separate software process in Junos OS and is not in the packet-forwarding path.

### Family mpls

You must configure the core interfaces used for MPLS traffic to belong to **family mpls**.



**NOTE:** You can enable **family mpls** on either individual interfaces or on aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

#### Related Documentation

- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)
- [Understanding Using MPLS-Based Layer 3 VPNs on Switches on page 162](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 158](#)
- [Configuring MPLS on Provider Edge Switches on page 188](#)
- [Configuring MPLS on Provider Switches on page 192](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 199](#)
- [Configuring a Global MPLS EXP Classifier on page 184](#)
- [Configuring Ethernet over MPLS \(L2 Circuit\) on page 180](#)
- [Junos OS MPLS Applications Library for Routing Devices](#)
- [Junos OS VPNs Library for Routing Devices](#)

## Understanding MPLS Label Operations

---

In the traditional packet-forwarding paradigm, as a packet travels from one switch to the next, an independent forwarding decision is made at each hop. The IP network 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 made only once, when a packet enters the MPLS tunnel (that is, the path used for MPLS traffic).

When an IP packet enters a label-switched path (LSP), the ingress provider edge (PE) switch 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 provider switch in the LSP. This switch and all subsequent switches 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 switch in the path. When the packet reaches the egress PE switch, the label is removed, and the packet again becomes a native IP packet and is forwarded based on its IP routing information.

This topic describes:


- [MPLS Label-Switched Paths and MPLS Labels on page 150](#)
- [Reserved Labels on page 151](#)
- [MPLS Label Operations on page 151](#)
- [Penultimate-Hop Popping and Ultimate-Hop Popping on page 153](#)

### MPLS Label-Switched Paths and MPLS Labels

When a packet enters the MPLS network, it is assigned to an LSP. Each LSP is identified by a label, which is a short (20-bit), fixed-length value at the front of the MPLS label (32 bits). Labels are used as lookup indexes for the label forwarding table. For each label, this table stores forwarding information. Because no additional parsing or lookup is done on the encapsulated packet, MPLS supports the transmission of any other protocols within the packet payload.

[Figure 3 on page 151](#) shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.

Figure 3: Label Encoding



## Reserved Labels

Labels range from 0 through 1,048,575. Labels 0 through 999,999 are for internal use.

Some of the reserved labels (in the range 0 through 15) have well-defined meanings.

The following reserved labels are used by QFX Series and EX4600 devices:

- 0, IPv4 Explicit Null label—This value is valid 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 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.
- 3, Implicit Null label—This label is used in the signaling protocol (RSVP) only to request label popping by the downstream switch. It never actually appears in the encapsulation. Labels with a value of 3 must not be used in the data packet as real labels. No payload type (IPv4 or IPv6) is implied with this label.

## MPLS Label Operations

QFX Series and EX4600 devices support the following MPLS label operations:

- Push
- Pop
- Swap



**NOTE:** There is a limit with regard to the number of labels that QFX and EX4600 devices can affix (push operations) to the label stack or remove (pop operations) from the label stack.

- For Push operations—As many as three labels are supported.
- For Pop operations—As many as three labels are supported.

The push operation affixes a new label to the top of the IP packet. For IPv4 packets, the new label is the first label. The time to live (TTL) field value in the packet header is derived from the IP packet header. The push operation cannot be applied to a packet that already has an MPLS label.

The pop operation removes a 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.

The swap operation removes an existing MPLS label from an IP packet and replaces it with a new MPLS label, based on the following:

- Incoming interface
- Label
- Label forwarding table

Figure 4 on page 152 shows an IP packet without a label arriving on the customer edge interface (ge-0/0/1) of the ingress PE switch. The ingress PE switch examines the packet and identifies that packet's destination as the egress PE switch. The ingress PE switch applies label 100 to the packet and sends the MPLS packet to its outgoing MPLS core interface (ge-0/0/5). The MPLS packet is transmitted on the MPLS tunnel through the provider switch, where it arrives at interface ge-0/0/5 with label 100. The provider switch swaps label 100 with label 200 and forwards the MPLS packet through its core interface (ge-0/0/7) to the next hop on the tunnel, which is the egress PE switch. The egress PE switch receives the MPLS packet through its core interface (ge-0/0/7), removes the MPLS label, and sends the IP packet out of its customer edge interface (ge-0/0/1) to a destination that is beyond the tunnel.

**Figure 4: MPLS Label Swapping**




Figure 4 on page 152 shows the path of a packet as it passes in one direction from the ingress PE switch to the egress PE switch. However, the MPLS configuration also allows traffic to travel in the reverse direction. Thus, each PE switch operates as both an ingress switch and an egress switch.



## Penultimate-Hop Popping and Ultimate-Hop Popping

The switches enable penultimate-hop popping (PHP) by default with IP over MPLS configurations. With PHP, the penultimate provider switch is responsible for popping the MPLS label and forwarding the traffic to the egress PE switch. The egress PE switch then performs an IP route lookup and forwards the traffic. This reduces the processing load on the egress PE switch, because it is not responsible for popping the MPLS label.

- The default advertised label is label 3 (Implicit Null label). If label 3 is advertised, the penultimate-hop switch removes the label and sends the packet to the egress PE switch.
- If ultimate-hop popping is enabled, label 0 (IPv4 Explicit Null label) is advertised and the egress PE switch of the LSP removes the label.

### **Related Documentation**

- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 146](#)
- [Configuring MPLS on Provider Edge Switches on page 188](#)
- [Configuring MPLS on Provider Switches on page 192](#)
- *Junos OS MPLS Applications Library for Routing Devices*
- *Junos OS VPNs Library for Routing Devices*

## Understanding BGP

---

BGP is an exterior gateway protocol (EGP) that is used to exchange routing information among routers in different autonomous systems (ASs). BGP routing information includes the complete route to each destination. BGP uses the routing information to maintain a database of network reachability information, which it exchanges with other BGP systems. BGP uses the network reachability information to construct a graph of AS connectivity, which enables BGP to remove routing loops and enforce policy decisions at the AS level.

Multiprotocol BGP (MBGP) extensions enable BGP to support IP version 6 (IPv6). MBGP defines the attributes `MP_REACH_NLRI` and `MP_UNREACH_NLRI`, which are used to carry IPv6 reachability information. Network layer reachability information (NLRI) update messages carry IPv6 address prefixes of feasible routes.

BGP allows for policy-based routing. You can use routing policies to choose among multiple paths to a destination and to control the redistribution of routing information.

BGP uses TCP as its transport protocol, using port 179 for establishing connections. Running over a reliable transport protocol eliminates the need for BGP to implement update fragmentation, retransmission, acknowledgment, and sequencing.

The Junos OS routing protocol software supports BGP version 4. This version of BGP adds support for Classless Interdomain Routing (CIDR), which eliminates the concept of network classes. Instead of assuming which bits of an address represent the network by looking at the first octet, CIDR allows you to explicitly specify the number of bits in the network address, thus providing a means to decrease the size of the routing tables. BGP version 4 also supports aggregation of routes, including the aggregation of AS paths.

This section discusses the following topics:

- [Autonomous Systems on page 154](#)
- [AS Paths and Attributes on page 154](#)
- [External and Internal BGP on page 155](#)
- [Multiple Instances of BGP on page 155](#)

### Autonomous Systems

An *autonomous system* (AS) is a set of routers that are under a single technical administration and normally use a single interior gateway protocol and a common set of metrics to propagate routing information within the set of routers. To other ASs, an AS appears to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it.

### AS Paths and Attributes

The routing information that BGP systems exchange includes the complete route to each destination, as well as additional information about the route. The route to each destination is called the *AS path*, and the additional route information is included in *path attributes*. BGP uses the AS path and the path attributes to completely determine the network topology. Once BGP understands the topology, it can detect and eliminate


routing loops and select among groups of routes to enforce administrative preferences and routing policy decisions.

## External and Internal BGP

BGP supports two types of exchanges of routing information: exchanges among different ASs and exchanges within a single AS. When used among ASs, BGP is called *external BGP* (EBGP) and BGP sessions perform *inter-AS routing*. When used within an AS, BGP is called *internal BGP* (IBGP) and BGP sessions perform *intra-AS routing*.

Figure 5 on page 155 illustrates ASs, IBGP, and EBGP.

Figure 5: ASs, EBGP, and IBGP



A BGP system shares network reachability information with adjacent BGP systems, which are referred to as *neighbors* or *peers*.

BGP systems are arranged into *groups*. In an IBGP group, all peers in the group—called *internal peers*—are in the same AS. Internal peers can be anywhere in the local AS and do not have to be directly connected to one another. Internal groups use routes from an IGP to resolve forwarding addresses. They also propagate external routes among all other internal routers running IBGP, computing the next hop by taking the BGP next hop received with the route and resolving it using information from one of the interior gateway protocols.

In an EBGP group, the peers in the group—called *external peers*—are in different ASs and normally share a subnet. In an external group, the next hop is computed with respect to the interface that is shared between the external peer and the local router.

## Multiple Instances of BGP

You can configure multiple instances of BGP at the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols]

Multiple instances of BGP are primarily used for Layer 3 VPN support.

IGP peers and external BGP (EBGP) peers (both nonmultihop and multihop) are all supported for routing instances. BGP peering is established over one of the interfaces configured under the **routing-instances** hierarchy.



**NOTE:** When a BGP neighbor sends BGP messages to the local routing device, the incoming interface on which these messages are received must be configured in the same routing instance that the BGP neighbor configuration exists in. This is true for neighbors that are a single hop away or multiple hops away.

Routes learned from the BGP peer are added to the **instance-name.inet.0** table by default. You can configure import and export policies to control the flow of information into and out of the instance routing table.

For Layer 3 VPN support, configure BGP on the provider edge (PE) router to receive routes from the customer edge (CE) router and to send the instances' routes to the CE router if necessary. You can use multiple instances of BGP to maintain separate per-site forwarding tables for keeping VPN traffic separate on the PE router.

You can configure import and export policies that allow the service provider to control and rate-limit traffic to and from the customer.

You can configure an EBGP multihop session for a VRF routing instance. Also, you can set up the EBGP peer between the PE and CE routers by using the loopback address of the CE router instead of the interface addresses.

**Related Documentation**

- *BGP Routes Overview*
- *BGP Messages Overview*

---

## IPv6 Layer 3 VPNs

The interfaces between the PE and CE routers of a Layer 3 VPN can be configured to carry IP version 6 (IPv6) traffic. IP allows numerous nodes on different networks to interoperate seamlessly. IPv4 is currently used in intranets and private networks, as well as the Internet. IPv6 is the successor to IPv4, and is based for the most part on IPv4.

In the Juniper Networks implementation of IPv6, the service provider implements an MPLS-enabled IPv4 backbone to provide VPN service for IPv6 customers. The PE routers have both IPv4 and IPv6 capabilities. They maintain IPv6 VPN routing and forwarding (VRF) tables for their IPv6 sites and encapsulate IPv6 traffic in MPLS frames that are then sent into the MPLS core network.

IPv6 for Layer 3 VPNs is supported for BGP and for static routes.

IPv6 over Layer 3 VPNs is described in RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*.

For more information about IPv6, see the *Junos OS Routing Protocols Library for Routing Devices*.

## Ethernet Pseudowire Overview

---

An Ethernet pseudowire is used to carry Ethernet or 802.3 Protocol Data Units (PDUs) over an MPLS network enabling service providers to offer emulated Ethernet services over existing MPLS networks. Ethernet or 802.3 PDUs are encapsulated within the pseudowire to provide a point-to-point Ethernet service. For the point-to-point Ethernet service, the following fault management features are supported:

- The IEEE 802.3ah standard for Operation, Administration, and Management (OAM). You can configure IEEE 802.3ah OAM link-fault management on Ethernet point-to-point direct links or links across Ethernet repeaters.

Ethernet OAM link-fault management can be used for physical link-level fault detection and management. It uses a new, optional sublayer in the data link layer of the OSI model. Ethernet OAM can be implemented on any full-duplex point-to-point or emulated point-to-point Ethernet link. A system-wide implementation is not required; OAM can be deployed on particular interfaces of a router. Transmitted Ethernet OAM messages or OAM PDUs are of standard length, untagged Ethernet frames within the normal frame length limits in the range 64–1518 bytes.

- Ethernet connectivity fault management (CFM) to monitor the physical link between two routers.
  - Connection protection using the continuity check protocol for fault monitoring . The continuity check protocol is a neighbor discovery and health check protocol that discovers and maintains adjacencies at the VLAN or link level.
  - Path protection using the linktrace protocol for path discovery and fault verification . Similar to IP traceroute, the linktrace protocol maps the path taken to a destination MAC address through one or more bridged networks between the source and destination.

### Related Documentation

- *Configuring IEEE 802.3ah OAM Link-Fault Management*
- *Pseudowire Overview for ACX Series Universal Access Routers*
- *TDM Pseudowires Overview*
- *ATM Pseudowire Overview*

## Understanding CoS MPLS EXP Classifiers and Rewrite Rules

---

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion by applying packet classifiers and rewrite rules to the MPLS traffic. MPLS classifiers are global and apply to all interfaces configured as **family mpls** interfaces.

When a packet enters a customer-edge interface on the ingress provider edge (PE) switch, the switch associates the packet with a particular CoS servicing level before placing the packet onto the label-switched path (LSP). The switches within the LSP utilize the CoS value set at the ingress PE switch to determine the CoS service level. The CoS value embedded in the classifier is translated and encoded in the MPLS header by means of the experimental (EXP) bits.

EXP classifiers map incoming MPLS packets to a forwarding class and a loss priority, and assign MPLS packets to output queues based on the forwarding class mapping. EXP classifiers are behavior aggregate (BA) classifiers.

EXP rewrite rules change (rewrite) the CoS value of the EXP bits in outgoing packets on the egress queues of the switch so that the new (rewritten) value matches the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.



**NOTE:** On QFX5200, QFX5100, QFX3500, QF3600, and EX4600 switches, and on QFabric systems, there is no default EXP classifier. If you want to classify incoming MPLS packets using the EXP bits, you must configure a global EXP classifier. The global EXP classifier applies to all MPLS traffic on interfaces configured as **family mpls**.

On QFX10000 switches, there is a no default EXP classifier. If you want to classify incoming MPLS packets using the EXP bits, you must configure EXP classifiers and apply them to logical interfaces configured as **family mpls**. (You cannot apply classifiers to physical interfaces.). You can configure up to 64 EXP classifiers.

There is no default EXP rewrite rule. If you want to rewrite the EXP bit value at the egress interface, you must configure EXP rewrite rules and apply them to logical interfaces.

EXP classifiers and rewrite rules are applied only to interfaces that are configured as **family mpls** (for example, set interfaces **xe-0/0/35 unit 0 family mpls.**)

---

This topic includes:

- [EXP Classifiers on page 159](#)
- [EXP Rewrite Rules on page 160](#)
- [Schedulers on page 161](#)

## EXP Classifiers

On QFX5200, QFX5100, EX4600, QFX3500, and QFX3600 switches, and on QFabric systems, unlike DSCP and IEEE 802.1p BA classifiers, EXP classifiers are global to the switch and apply to all switch interfaces that are configured as **family mpls**. On QFX10000 switches, you apply EXP classifiers to individual logical interfaces, and different interfaces can use different EXP classifiers.

When you configure and apply an EXP classifier, MPLS traffic on all **family mpls** interfaces uses the EXP classifier, even on interfaces that also have a fixed classifier. If an interface has both an EXP classifier and a fixed classifier, the EXP classifier is applied to MPLS traffic and the fixed classifier is applied to all other traffic.

Also unlike DSCP and IEEE 802.1p BA classifiers, there is no default EXP classifier. If you want to classify MPLS traffic based on the EXP bits, you must explicitly configure an EXP classifier and apply it to the switch interfaces. Each EXP classifier has eight entries that correspond to the eight EXP CoS values (0 through 7, which correspond to CoS bits 000 through 111).

You can configure up to 64 EXP classifiers.

However, on QFX5200, QFX5100, EX4600, and legacy CLI switches, the switch uses only one MPLS EXP classifier as a global classifier on all interfaces. After you configure an MPLS EXP classifier, you can configure that classifier as the global EXP classifier by including the EXP classifier in the **[edit class-of-service system-defaults classifiers exp]** hierarchy level. All switch interfaces configured as **family mpls** use the global EXP classifier to classify MPLS traffic.

On these switches, only one EXP classifier can be configured as the global EXP classifier at any time. If you want to change the global EXP classifier, delete the global EXP classifier configuration (use the **user@switch# delete class-of-service system-defaults classifiers exp** configuration statement), then configure the new global EXP classifier.

QFX10000 switches do not support global EXP classifiers. You can configure one EXP classifier and apply it to multiple logical interfaces, or configure multiple EXP classifiers and apply different EXP classifiers to different logical interfaces.

If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. (Switches that have a default EXP classifier use the default classifier.) If no EXP classifier and no fixed classifier are applied to the interface, MPLS traffic is treated as best-effort traffic using the 802.1 default untrusted classifier. DSCP classifiers are not applied to MPLS traffic.

On QFX5200, QFX5100, EX4600, and legacy CLI switches, because the EXP classifier is global, you cannot configure some ports to use a fixed IEEE 802.1p classifier for MPLS traffic on some interfaces and the global EXP classifier for MPLS traffic on other interfaces. When you configure a global EXP classifier, all MPLS traffic on all interfaces uses the EXP classifier.



NOTE: The switch uses only the outermost label of incoming EXP packets for classification.



NOTE: MPLS packets with 802.1Q tags are not supported.

## EXP Rewrite Rules

As MPLS packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. EXP rewrite rules set the value of the EXP CoS bits within the header of the outgoing MPLS packet on **family mpls** interfaces. Each rewrite rule reads the current forwarding class and loss priority associated with the packet, locates the chosen CoS value from a table, and writes that CoS value into the packet header, replacing the old CoS value. EXP rewrite rules apply only to MPLS traffic.

EXP rewrite rules apply only to logical interfaces. You cannot apply EXP rewrite rules to physical interfaces.

There are no default EXP rewrite rules. If you want to rewrite the EXP value in MPLS packets, you must configure EXP rewrite rules and apply them to logical interfaces. If no rewrite rules are applied, all MPLS labels that are pushed have a value of zero (0). The EXP value remains unchanged on MPLS labels that are swapped.

You can configure up to 64 EXP rewrite rules, but you can only apply 16 EXP rewrite rules at any time on the switch. On a given logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.

You can apply an EXP rewrite rule to an interface that has a DSCP, DSCP IPv6, or IEEE 802.1p rewrite rule. Only MPLS traffic uses the EXP rewrite rule. MPLS traffic does not use DSCP or DSCP IPv6 rewrite rules.

If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.



NOTE: On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.



## Schedulers

The schedulers for using CoS with MPLS are the same as for the other CoS configurations on the switch. Default schedulers are provided only for the best-effort, fcoe, no-loss, and network-control default forwarding classes. If you configure a custom forwarding class for MPLS traffic, you need to configure a scheduler to support that forwarding class and provide bandwidth to that forwarding class.

### Related Documentation

- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces](#)

## Understanding Ethernet-over-MPLS (L2 Circuit)

Ethernet-over-MPLS allows sending Layer 2 (L2) Ethernet frames transparently over MPLS. Ethernet-over-MPLS uses a tunneling mechanism for Ethernet traffic through an MPLS-enabled Layer 3 core. It encapsulates Ethernet protocol data units (PDUs) inside MPLS packets and forwards the packets, using label stacking, across the MPLS network. This technology has applications in service provider, enterprise and data center environments. For disaster recovery purposes, data centers are hosted in multiple sites that are geographically distant and interconnected using a WAN network.



**NOTE:** A Layer 2 circuit is similar to a circuit cross-connect (CCC), except that multiple Layer 2 circuits can be transported over a single label-switched path (LSP) tunnel between two provider edge (PE) routers. In contrast, each CCC requires a dedicated LSP.

- [Ethernet-over-MPLS in Data Centers on page 161](#)

## Ethernet-over-MPLS in Data Centers

For disaster recovery purposes, data centers are hosted in multiple sites that are geographically distant and interconnected using a WAN network. These data centers require L2 connectivity between them for the following reasons:

- To replicate the storage over Fiber Channel IP (FCIP). FCIP works only on the same broadcast domain.
- To run a dynamic routing protocol between the sites.
- To support High Availability clusters that interconnect the nodes hosted in the various data centers.

### Related Documentation

- [Configuring Ethernet over MPLS \(L2 Circuit\) on page 180](#)

## Understanding Using MPLS-Based Layer 3 VPNs on Switches

---

On the QFX Series switches and on EX4600 switches, you can use MPLS-based Layer 3 virtual private networks (VPNs) to securely connect geographically diverse sites across an MPLS network. MPLS services can be used to connect various sites to a backbone network and to ensure better performance for low-latency applications such as voice over IP (VoIP) and other business-critical functions.

A VPN uses a public telecommunications infrastructure, such as the Internet, to provide remote offices or individual users with secure access to their organization's network. VPNs are designed to provide the same level of performance and security as privately owned or leased networks but without the attendant costs.

This topic describes:

- [MPLS-Based Layer 3 VPNs on page 162](#)

### MPLS-Based Layer 3 VPNs

In Junos OS, Layer 3 VPNs are based on RFC 4364, *BGP/MPLS IP Virtual Private Networks*. RFC 4364 defines a mechanism by which service providers can use their IP backbones to provide VPN services to their customers. A Layer 3 VPN is a set of sites that share common routing information and whose connectivity is controlled by a collection of policies. The sites that make up a Layer 3 VPN are connected over a provider's existing public Internet backbone.

Customer networks, because they are private, can use either public or private addresses, as defined in RFC 1918, *Address Allocation for Private Internets*. When customer networks that use private addresses connect to the public Internet infrastructure, the private addresses might overlap with the same private addresses used by other network users. BGP/MPLS VPNs solve this problem by adding a VPN identifier prefix to each address from a particular VPN site, thereby creating an address that is unique both within the VPN and on the public Internet. In addition, each VPN has its own VPN-specific routing table that contains the routing information for that VPN only. Two different VPNs can use overlapping addresses. Each route within a VPN is assigned an MPLS label (for example, MPLS-ARCH, MPLS-BGP, or MPLS-ENCAPS). When BGP distributes a VPN route, it also distributes an MPLS label for that route. Before a customer data packet travels across the service provider's backbone, it is encapsulated along with the MPLS label that corresponds to the route within the customer's VPN that is the best match based on the packet's destination address. This MPLS packet is further encapsulated with another MPLS label or with an IP, so that it gets tunneled across the backbone to the egress provider edge (PE) switch. Thus, the backbone core switches do not need to know the VPN routes.

QFX5100 switches also support interprovider VPNs, and carrier-of-carriers VPNs. For more information, see ["Interprovider and Carrier-of-Carriers VPNs" on page 165](#)

#### Related Documentation

- [Understanding MPLS Label Operations on page 150](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 146](#)

- *Example: Configuring MPLS-Based Layer 2 VPNs*
- *Example: Configuring MPLS-Based Layer 3 VPNs on EX Series Switches*


## Carrier-of-Carriers VPNs

The customer of a VPN service provider might be a service provider for the end customer. The following are the two main types of carrier-of-carriers VPNs (as described in RFC 4364):

- [“Internet Service Provider as the Customer” on page 165](#)—The VPN customer is an ISP that uses the VPN service provider’s network to connect its geographically disparate regional networks. The customer does not have to configure MPLS within its regional networks.
- [“VPN Service Provider as the Customer” on page 165](#)—The VPN customer is itself a VPN service provider offering VPN service to its customers. The carrier-of-carriers VPN service customer relies on the backbone VPN service provider for inter-site connectivity. The customer VPN service provider is required to run MPLS within its regional networks.

[Figure 6 on page 164](#) illustrates the network architecture used for a carrier-of-carriers VPN service.

**Figure 6: Carrier-of-Carriers VPN Architecture**



807197

This topic covers the following:

- [Internet Service Provider as the Customer on page 165](#)
- [VPN Service Provider as the Customer on page 165](#)

## Internet Service Provider as the Customer

In this type of carrier-of-carriers VPN configuration, ISP A configures its network to provide Internet service to ISP B. ISP B provides the connection to the customer wanting Internet service, but the actual Internet service is provided by ISP A.

This type of carrier-of-carriers VPN configuration has the following characteristics:

- The carrier-of-carriers VPN service customer (ISP B) does not need to configure MPLS on its network.
- The carrier-of-carriers VPN service provider (ISP A) must configure MPLS on its network.
- MPLS must also be configured on the CE routers and PE routers connected together in the carrier-of-carriers VPN service customer's and carrier-of-carriers VPN service provider's networks.

## VPN Service Provider as the Customer

A VPN service provider can have customers that are themselves VPN service providers. In this type of configuration, also called a hierarchical or recursive VPN, the customer VPN service provider's VPN-IPv4 routes are considered external routes, and the backbone VPN service provider does not import them into its VRF table. The backbone VPN service provider imports only the customer VPN service provider's internal routes into its VRF table.

The similarities and differences between interprovider and carrier-of-carriers VPNs are shown in [Table 18 on page 165](#).

**Table 18: Comparison of Interprovider and Carrier-of-Carriers VPNs**

Feature	ISP Customer	VPN Service Provider Customer
Customer edge device	AS border router	PE router
IBGP sessions	Carry IPv4 routes	Carry external VPN-IPv4 routes with associated labels
Forwarding within the customer network	MPLS is optional	MPLS is required

## Interprovider and Carrier-of-Carriers VPNs

All interprovider and carrier-of-carriers VPNs share the following characteristics:

- Each interprovider or carrier-of-carriers VPN customer must distinguish between internal and external customer routes.

- Internal customer routes must be maintained by the VPN service provider in its PE routers.
- External customer routes are carried only by the customer's routing platforms, not by the VPN service provider's routing platforms.

The key difference between interprovider and carrier-of-carriers VPNs is whether the customer sites belong to the same AS or to separate ASs:

- *Interprovider VPNs*—The customer sites belong to different ASs. You need to configure EBGP to exchange the customer's external routes.
- ["Carrier-of-Carriers VPNs" on page 164](#)—The customer sites belong to the same AS. You need to configure IBGP to exchange the customer's external routes.

In general, each service provider in a VPN hierarchy is required to maintain its own internal routes in its P routers, and the internal routes of its customers in its PE routers. By recursively applying this rule, it is possible to create a hierarchy of VPNs.

The following are definitions of the types of PE routers specific to interprovider and carrier-of-carriers VPNs:

- The AS border router is located at the AS border and handles traffic leaving and entering the AS.
- The end PE router is the PE router in the customer VPN; it is connected to the CE router at the end customer's site.

**Related  
Documentation**

- [Carrier-of-Carriers VPNs on page 164](#)
- *Interprovider VPNs*

---

## Chained Composite Next Hops for Transit Devices for VPNs

The Juniper Networks PTX Series Packet Transport Routers, MX Series 3D Universal Edge Routers with MIC and MPC interfaces, T4000 Core Routers, and QFX10000 switches are principally designed to handle large volumes of transit traffic in the core of large networks. Chained composite next hops help to facilitate this capability by allowing the router to process much larger volumes of routes. A chained composite next hop allows the router to direct sets of routes sharing the same destination to a common forwarding next hop, rather than having each route also include the destination. In the event that a network destination is changed, rather than having to update all of the routes sharing that destination with the new information, just the shared forwarding next hop is updated with the new information. The chained composite next hops continue to point to this forwarding next hop which now contains the new destination.

When the next hops for MPLS LSPs are created on the routers, the tag information corresponding to the inner-most MPLS label is extracted into a chained composite next hop. The chained composite next hop is stored in the ingress PFE. The chained composite next hop points to a next hop called the forwarding next hop that resides on the egress PFE. The forwarding next hop contains all of the other information (all of the labels

except for the inner-most labels; and the IFA/IP information corresponding to the actual next hop node). Many chained composite next hops can share the same forwarding next hop. Additionally, separating the label from the forwarding next hop and storing it on the ingress PFE (within the chained composite next hop) helps to conserve egress PFE memory by reducing the number of rewrite strings stored on the egress PFE.

The support of chained composite next hops for directly connected Provider Edge (PE) routers varies from one platform to another.

On platforms containing only MPCs, such as PTX Series Packet Transport Routers, the MX80 router, the MX2020 router, and the QFX10000 switches, chained composite next hops are enabled by default for the following MPLS and VPN protocols and applications:



**NOTE:** Point-to-Multipoint LSPs and Layer 2 VPNs are not supported on the QFX10000 switches.

- Labeled BGP
- Layer 2 VPNs
- Layer 3 VPNs
- LDP
- MPLS
- Point-to-Multipoint LSPs
- RSVP
- Static LSPs

On MX Series 3D Universal Edge Routers containing both DPC and MPC FPCs and on T4000 Core Routers containing MPC and FPCs, chained composite next hops are disabled by default and need to be explicitly configured.

To enable chained composite next hops include the **l3vpn** statement at the **[edit routing-options forwarding-table chained-composite-next-hop ingress]** hierarchy level.

Starting with Junos OS Release 13.3, for chained composite next hop feature to take effect for directly connected PE devices, the chassis must be configured to use the **enhanced-ip** option (in the case of MX Series 3D Universal Edge Routers containing both DPC and MPC FPCs) or the **enhanced-mode** option (in the case of T4000 Core Routers containing MPC and FPCs) in the network service mode, in addition to the **l3vpn** configuration.

For more information about configuring chassis network services, see the *Junos OS Administration Library for Routing Devices*.

#### Related Documentation

- *Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs*
- *Example: Configuring Chained Composite Next Hops for Direct PE-PE Connections in VPNs*


- [transit \(Chained Composite Next Hops\) on page 306](#)

## 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 7 on page 168](#) illustrates an LSP from Router A to Router F, showing the established detours. Each detour is established by an upstream node to avoid the link toward the immediate downstream node and the immediate downstream node itself. Each detour might traverse through one or more label-switched routers (or switches) that are not shown in the figure.

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


**Figure 7: 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 8 on page 168](#) illustrates the detour taken when the link between Router B and Router C fails.

**Figure 8: 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 7 on page 168](#) 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 9 on page 169](#), 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 9: Detours Merging into Other Detours**



g017074

**Related  
Documentation**

- [fast-reroute on page 247](#)
- [Configuring Fast Reroute](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)
- [Interprovider and Carrier-of-Carriers VPNs on page 165](#)

---

## Graceful Restart and MPLS-Related Protocols

---

This section contains the following topics:

- [LDP on page 170](#)
- [RSVP on page 171](#)
- [CCC and TCC on page 171](#)

### LDP

LDP graceful restart enables a router whose LDP control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers. It also enables a router on which helper mode is enabled to assist a neighboring router that is attempting to restart LDP.

During session initialization, a router advertises its ability to perform LDP graceful restart or to take advantage of a neighbor performing LDP graceful restart by sending the graceful restart TLV. This TLV contains two fields relevant to LDP graceful restart: the reconnect time and the recovery time. The values of the reconnect and recovery times indicate the graceful restart capabilities supported by the router.

The reconnect time is configured in Junos OS as 60 seconds and is not user-configurable. The reconnect time is how long the helper router waits for the restarting router to establish a connection. If the connection is not established within the reconnect interval, graceful restart for the LDP session is terminated. The maximum reconnect time is 120 seconds and is not user-configurable. The maximum reconnect time is the maximum value that a helper router accepts from its restarting neighbor.

When a router discovers that a neighboring router is restarting, it waits until the end of the recovery time before attempting to reconnect. The recovery time is the length of time a router waits for LDP to restart gracefully. The recovery time period begins when an initialization message is sent or received. This time period is also typically the length of time that a neighboring router maintains its information about the restarting router, so it can continue to forward traffic.

You can configure LDP graceful restart both in the master instance for the LDP protocol and for a specific routing instance. You can disable graceful restart at the global level for all protocols, at the protocol level for LDP only, and for a specific routing instance only.

## RSVP

RSVP graceful restart enables a router undergoing a restart to inform its adjacent neighbors of its condition. The restarting router requests a grace period from the neighbor or peer, which can then cooperate with the restarting router. The restarting router can still forward MPLS traffic during the restart period; convergence in the network is not disrupted. The restart is not visible to the rest of the network, and the restarting router is not removed from the network topology. RSVP graceful restart can be enabled on both transit routers and ingress routers. It is available for both point-to-point LSPs and point-to-multipoint LSPs.

## CCC and TCC

CCC and TCC graceful restart enables Layer 2 connections between customer edge (CE) routers to restart gracefully. These Layer 2 connections are configured with the **remote-interface-switch** or **lsp-switch** statements. Because these CCC and TCC connections have an implicit dependency on RSVP LSPs, graceful restart for CCC and TCC uses the RSVP graceful restart capabilities.

RSVP graceful restart must be enabled on the provider edge (PE) routers and provider (P) routers to enable graceful restart for CCC and TCC. Also, because RSVP is used as the signaling protocol for signaling label information, the neighboring router must use helper mode to assist with the RSVP restart procedures.

### Related Documentation

- *Graceful Restart Concepts*
- *Graceful Restart System Requirements*
- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring Graceful Restart*

## Types of LSPs

---

There are three types of LSPs:

- Static LSPs—For static paths, you must manually assign labels on all routers involved (ingress, transit, and egress). No signaling protocol is needed. This procedure is similar to configuring static routes on individual routers. Like static routes, there is no error reporting, liveliness detection, or statistics reporting.
- LDP-signaled LSPs—See [“LDP Introduction” on page 4](#).
- 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.

## Configuring Automatic Bandwidth Allocation for LSPs

---

Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure an LSP with minimal bandwidth, and this feature can dynamically adjust the LSP's bandwidth allocation based on current traffic patterns. The bandwidth adjustments do not interrupt traffic flow through the tunnel.

At the end of the automatic bandwidth allocation time interval, the current maximum average bandwidth usage is compared with the allocated bandwidth for the LSP. If the LSP needs more bandwidth, an attempt is made to set up a new path where bandwidth is equal to the current maximum average usage. If the attempt is successful, the LSP's traffic is routed through the new path and the old path is removed. If the attempt fails, the LSP continues to use its current path.

If you have configured link and node protection for the LSP and traffic has been switched to the bypass LSP, the automatic bandwidth allocation feature continues to operate and take bandwidth samples from the bypass LSP. For the first bandwidth adjustment cycle, the maximum average bandwidth usage taken from the original link and node-protected LSP is used to resignal the bypass LSP if more bandwidth is needed. (Link and node protection are not supported on QFX Series switches.)

If you have configured fast-reroute for the LSP, you might not be able to use this feature to adjust the bandwidth. Because the LSPs use a fixed filter (FF) reservation style, when a new path is signaled, the bandwidth might be double-counted. Double-counting can prevent a fast-reroute LSP from ever adjusting its bandwidth when automatic bandwidth allocation is enabled. (Fast reroute is not supported on QFX Series switches.)

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

- [Configuring Automatic Bandwidth Allocation on LSPs on page 173](#)
- [Requesting Automatic Bandwidth Allocation Adjustment on page 178](#)



**NOTE:** On the QFX10000 switches, you can only configure automatic bandwidth allocation at the `edit protocols mpls` hierarchy level. Logical systems are not supported.

## Configuring Automatic Bandwidth Allocation on LSPs

To enable automatic bandwidth allocation on an LSP, include the **auto-bandwidth** statement:

```
auto-bandwidth {
  adjust-interval seconds;
  adjust-threshold percent;
  adjust-threshold-overflow-limit number;
  adjust-threshold-underflow-limit number;
  maximum-bandwidth bps;
  minimum-bandwidth bps;
  minimum-bandwidth-adjust-interval
  minimum-bandwidth-adjust-threshold-change
  minimum-bandwidth-adjust-threshold-value
  monitor-bandwidth;
}
```

You can include this statement at the following hierarchy levels:

- `[edit protocols mpls label-switched-path lsp-name]`
- `[edit logical-systems logical-system-name protocols mpls label-switched-path lsp-name]`

If an LSP has an automatic bandwidth configuration, you can disable automatic bandwidth adjustments on a particular path (either primary or secondary) by configuring a static bandwidth value and by disabling the CSPF computation (using the **no-cspf** statement).

For example:

```
user@host> show protocols mpls
label-switched-path primary-path {
  to 192.168.0.1;
  ldp-tunneling;
  optimize-timer 3571;
  least-fill;
  link-protection;
  adaptive;
  auto-bandwidth {
    adjust-interval 7177;
    adjust-threshold 5;
    minimum-bandwidth 1m;
    maximum-bandwidth 2500000000;
    adjust-threshold-overflow-limit 2;
    resignal-minimum-bandwidth;
  }
  primary primary-path;
  secondary secondary-path {
    bandwidth 0;
    no-cspf;
```

```

        priority 0 0;
    }
}

```

The statements configured at the `[edit protocols mpls label-switched-path label-switched-path-name auto-bandwidth]` hierarchy level are optional and explained in the following sections:

- [Configuring the Automatic Bandwidth Allocation Interval on page 174](#)
- [Configuring the Maximum and Minimum Bounds of the LSP's Bandwidth on page 174](#)
- [Configuring the Automatic Bandwidth Adjustment Threshold on page 175](#)
- [Configuring a Limit on Bandwidth Overflow and Underflow Samples on page 175](#)
- [Configuring Passive Bandwidth Utilization Monitoring on page 177](#)

### Configuring the Automatic Bandwidth Allocation Interval

At the end of the automatic bandwidth allocation interval, the automatic bandwidth computation and new path setup process is triggered.



**NOTE:** To prevent unnecessary resignaling of LSPs, it is best to configure an LSP adjustment interval that is at least three times longer than the MPLS automatic bandwidth statistics interval. For example, if you configure a value of 30 seconds for the MPLS automatic bandwidth statistics interval (interval statement at the `[edit protocols mpls statistics]` hierarchy level), you should configure a value of at least 90 seconds for the LSP adjustment interval (adjust-interval statement at the `[edit protocols mpls label-switched-path label-switched-path-name auto-bandwidth]` hierarchy level). See also “[Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs](#)” on page 193.

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 resigaled with the new bandwidth.

During the automatic bandwidth adjustment interval, the router might receive a steady increase in traffic (increasing bandwidth utilization) on an LSP, potentially causing congestion or packet loss. To prevent this, you can define a second trigger to prematurely expire the automatic bandwidth adjustment timer before the end of the current adjustment interval.

Every statistics interval, the router samples the average bandwidth utilization of an LSP and if this has exceeded the current maximum average bandwidth utilization, the maximum average bandwidth utilization is updated.

During each sample period, the following conditions are also checked:

- Is the current average bandwidth utilization above the active bandwidth of the path?
- Has the difference between the average bandwidth utilization and the active bandwidth exceeded the adjust threshold (bandwidth utilization has changed significantly) ?

If these conditions are true, it is considered to be one bandwidth overflow sample. Using the **adjust-threshold-overflow-limit** statement, you can define a limit on the number of bandwidth overflow samples such that when the limit is reached, the current automatic bandwidth adjustment timer is expired and a bandwidth adjustment is triggered. Once this adjustment is complete, the normal automatic bandwidth adjustment timer is reset to expire after the periodic adjustment interval.

To specify a limit on the number of bandwidth overflow samples before triggering an automatic bandwidth allocation adjustment, configure the **adjust-threshold-overflow-limit** statement:

```
adjust-threshold-overflow-limit number;
```

Similarly, if the current average bandwidth utilization is below the active bandwidth of the path by the configured adjusted threshold (meaning that bandwidth utilization has gone down significantly), the sample is considered to be an underflow sample. The adjusted (new signaling) bandwidth after an adjustment due to underflow is the maximum average bandwidth among the underflow samples. You can specify a limit on the number of bandwidth underflow samples before triggering an automatic bandwidth allocation adjustment by configuring the **adjust threshold-underflow-limit** statement:

```
adjust-threshold-underflow-limit number;
```

These statements can be configured at the following hierarchy levels:

- [edit protocols mpls **label-switched-path** *lsp-name* **auto-bandwidth**]
- [edit logical-systems *logical-system-name* protocols mpls **label-switched-path** *lsp-name* **auto-bandwidth**]

You must configure the **adjust-threshold** and **minimum-bandwidth** statements whenever you configure the **adjust-threshold-underflow-limit** statement. You must configure the **adjust-threshold** and **maximum-bandwidth** statements whenever you configure the **adjust-threshold-overflow-limit** statement



- You must configure a nonzero value for the **adjust-threshold** statement if you configure the **adjust-threshold-overflow-limit** or **adjust-threshold-underflow-limit** statement.
- Any bandwidth increase or decrease below the value configured for the **adjust-threshold** statement does not constitute an overflow or underflow condition.
- To prevent unlimited increases in LSP bandwidth (to limit overflow beyond a certain bandwidth), you must also configure the **maximum-bandwidth** statement when you configure the **adjust-threshold-overflow-limit** statement.

The following describes the other aspects of the **adjust-threshold-overflow-limit** statement:

- It only applies to bandwidth overflows. If the bandwidth is decreasing, the normal automatic bandwidth adjustment interval is used.
- It does not affect manually triggered automatic bandwidth adjustment.
- It applies to single-class DiffServ-TE LSPs.
- Because the **adjust-threshold-overflow-limit** statement can trigger a bandwidth adjustment, it cannot be enabled at the same time as the **monitor-bandwidth** statement (for information about that statement, see [“Configuring Passive Bandwidth Utilization Monitoring” on page 177](#)).
- 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 177](#)), and want to initiate an immediate bandwidth adjustment.

To use the **request mpls lsp adjust-autobandwidth** command, the following must be true:

- Automatic bandwidth allocation must be enabled on the LSP.
- The criteria required to trigger a bandwidth adjustment have been met (the difference between the adjust bandwidth and the current LSP path bandwidth is greater than the threshold limit).

A manually triggered bandwidth adjustment operates only on the active LSP path. Also, if you have enabled periodic automatic bandwidth adjustment, the periodic automatic bandwidth adjustment parameters (the adjustment interval and the maximum average bandwidth) are not reset after a manual adjustment.

For example, suppose the periodic adjust interval is 10 hours and there are currently 5 hours remaining before an automatic bandwidth adjustment is triggered. If you initiate a manual adjustment with the **request mpls lsp adjust-autobandwidth** command, the adjust timer is not reset and still has 5 hours remaining.

To manually trigger a bandwidth allocation adjustment, you need to use the **request mpls lsp adjust-autobandwidth** command. You can trigger the command for all affected LSPs on the router, or you can specify a particular LSP:

```
user@host> request mpls lsp adjust-autobandwidth
```

Once you execute this command, the automatic bandwidth adjustment validation process is triggered. If all the criteria for adjustment are met, the LSP's active path bandwidth is adjusted to the adjusted bandwidth value determined during the validation process.

#### Related Documentation

- [Configuring MPLS to Gather Statistics on page 187](#)
- [Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs on page 193](#)
- [request mpls lsp adjust-autobandwidth on page 332](#)
- [show mpls lsp on page 386](#)

## Configuring CoS Bits for an MPLS Network

When traffic enters a labeled-switch path (LSP) tunnel, the CoS bits in the MPLS header are set in one of two ways:

- The number of the output queue into which the packet was buffered and the packet loss priority (PLP) bit are written into the MPLS header and are used as the packet's CoS value. This behavior is the default, and no configuration is required. The *Class of Service Feature Guide for Routing Devices* explains the IP CoS values, and summarizes how the CoS bits are treated.
- You set a fixed CoS value on all packets entering the LSP tunnel. A fixed CoS value means that all packets entering the LSP receive the same class of service.

To set a fixed CoS value on all packets entering the LSP:

1. Specify a class of service value for the LSP:



**NOTE:** 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.

```
[edit protocols mpls]
user@switch# set class-of-service cos-value
```

#### Related Documentation

- [Understanding CoS Classifiers](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 158](#)
- [Configuring a Global MPLS EXP Classifier on page 184](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 199](#)
- [Defining CoS Rewrite Rules](#)

## Configuring Ethernet over MPLS (L2 Circuit)

To implement Ethernet over MPLS, you must configure a Layer 2 circuit on the provider edge (PE) switches. No special configuration is required on the customer edge (CE) switches. The provider switches require MPLS and LDP to be configured on the interfaces that will be receiving and transmitting MPLS packets.




**NOTE:** A Layer 2 circuit is similar to a circuit cross-connect (CCC), except that multiple Layer 2 circuits can be transported over a single label-switched path (LSP) tunnel between two PE switches. In contrast, each CCC requires a dedicated LSP.

This topic describes how to configure the PE switches to support Ethernet over MPLS. You must configure interfaces and protocols on both the local PE (PE1) and the remote PE (PE2) switches. The interface configuration varies depending upon whether the Layer 2 circuit is port-based or VLAN-based.

Figure 10 on page 180 shows an example of a Layer 2 circuit configuration.

**Figure 10: Ethernet over MPLS Layer 2 Circuit**



**NOTE:** This topic refers to the local PE switch as PE1 and the remote PE switch as PE2. It also uses interface names rather than variables to help clarify the connections between the switches. The loopback addresses of the switches are configured as follows:

- PE1: 1.1.1.1
- PE2: 4.4.4.4

- [Configuring the Local PE Switch for Port-Based Layer 2 Circuit \(Pseudo-wire\) on page 181](#)
- [Configuring the Remote PE Switch for Port-Based Layer 2 Circuit \(Pseudo-wire\) on page 181](#)
- [Configuring the Local PE Switch for VLAN-Based Layer 2 Circuit on page 182](#)
- [Configuring the Remote PE Switch for VLAN-Based Layer 2 Circuit on page 183](#)

## Configuring the Local PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire)



**CAUTION:** Configure MPLS networks with an MTU (maximum transmission unit) that is at least 12 bytes larger than the largest frame size that will be transported by the LSPs. If the size of an encapsulated packet on the ingress LSR exceeds the LSP MTU, that packet is dropped. If an egress LSR receives a packet on a VC LSP with a length (after the label stack and sequencing control word have been popped) that exceeds the MTU of the destination layer 2 interface, that packet is also dropped.

To configure the local PE switch (PE1) for a port-based layer 2 circuit (pseudo-wire):

1. Configure an access CE-facing interface for Ethernet encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation ethernet-ccc
```



**NOTE:** On QFX Series switches, the L2 circuit CE facing interface does not support Aggregated Ethernet (AE) interfaces.

2. Configure the Layer 2 circuit from PE1 to PE2:

```
[edit protocols]
user@switch# set l2circuit neighbor 4.4.4.4 interface xe-0/0/6 virtual-circuit-id 1
```

3. Configure the label switched path from PE1 to PE2:

```
[edit protocols]
user@switch# set mpls label-switched-path PE1-to-PE2 to 4.4.4.4
```

4. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch# set mpls interface xe-0/0/7
user@switch# set rsdp interface xe-0/0/7
user@switch# set ldp interface lo0.0
```

## Configuring the Remote PE Switch for Port-Based Layer 2 Circuit (Pseudo-wire)

To configure the remote PE switch (PE2) for a port-based layer 2 circuit:

1. Configure an access CE-facing interface for Ethernet encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation ethernet-ccc
```



**NOTE:** On QFX Series switches, the L2 circuit CE facing interface does not support AE interfaces.

2. Configure the Layer 2 circuit from PE2 to PE1:

```
[edit protocols]
user@switch# set l2circuit neighbor 1.1.1.1 interface xe-0/0/6 virtual-circuit-id 1
```

3. Configure the label switched path from PE2 to PE1:

- ```
[edit protocols]
user@switch#set mpls label-switched-path PE2-to-PE1 to 1.1.1.1
```
4. Configure the protocols on the core and loopback interfaces:
 

```
[edit protocols]
user@switch#set mpls interface xe-0/0/7
user@switch#set rsvp interface xe-0/0/7
user@switch#set ldp interface lo0.0
```

## Configuring the Local PE Switch for VLAN-Based Layer 2 Circuit

To configure the local PE switch (PE1) for a VLAN-based layer 2 circuit:

1. Configure an access CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation vlan-ccc
```



**NOTE:** On QFX Series switches, the L2 circuit CE facing interface does not support AE interfaces.

2. Configure the logical unit of the CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 encapsulation vlan-ccc
```

3. Configure the logical unit of the CE-facing interface to belong to family ccc:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 family ccc
```

4. Configure the same interface for VLAN tagging:

```
[edit interfaces]
user@switch# set xe-0/0/6 vlan-tagging
```

5. Configure the VLAN ID of the interface:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 vlan-id 600
```

6. Configure the Layer 2 circuit from PE1 to PE2:

```
[edit protocols]
user@switch#set l2circuit neighbor 4.4.4.4 interface xe-0/0/6 virtual-circuit-id 1
```

7. Configure the label switched path from PE1 to PE2:

```
[edit protocols]
user@switch#set mpls label-switched-path PE1-to-PE2 to 4.4.4.4
```

8. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch#set mpls interface xe-0/0/7
user@switch#set rsvp interface xe-0/0/7
user@switch#set ldp interface lo0.0
```

## Configuring the Remote PE Switch for VLAN-Based Layer 2 Circuit

To configure the remote PE switch (PE2) for a VLAN-based layer 2 circuit:

1. Configure an access CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 encapsulation vlan-ccc
```



**NOTE:** On QFX Series switches, the L2 circuit CE facing interface does not support AE interfaces.

2. Configure the logical unit of the CE-facing interface for VLAN encapsulation:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 encapsulation vlan-ccc
```

3. Configure the logical unit of the CE-facing interface to belong to family ccc:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 family ccc
```

4. Configure the same interface for VLAN tagging:

```
[edit interfaces]
user@switch# set xe-0/0/6 vlan-tagging
```

5. Configure the VLAN ID of the interface:

```
[edit interfaces]
user@switch# set xe-0/0/6 unit 0 vlan-id 600
```

6. Configure the Layer 2 circuit from PE2 to PE1:

```
[edit protocols]
user@switch# set l2circuit neighbor 1.1.1.1 interface xe-0/0/6 virtual-circuit-id 1
```

7. Configure the label switched path from PE2 to PE1:

```
[edit protocols]
user@switch# set mpls label-switched-path PE2-to-PE1 to 1.1.1.1
```

8. Configure the protocols on the core and loopback interfaces:

```
[edit protocols]
user@switch# set mpls interface xe-0/0/7
user@switch# set rsvp interface xe-0/0/7
user@switch# set ldp interface lo0.0
```

### Related Documentation

- [Understanding Ethernet-over-MPLS \(L2 Circuit\) on page 161](#)

## Configuring a Global MPLS EXP Classifier

---

EXP packet classification associates incoming packets with a particular MPLS CoS servicing level. EXP behavior aggregate (BA) classifiers examine the MPLS EXP value in the packet header to determine the CoS settings applied to the packet. EXP BA classifiers allow you to set the forwarding class and loss priority of an MPLS packet based on the incoming CoS value.

You can configure up to 64 EXP classifiers, however, the switch uses only one MPLS EXP classifier as a global classifier, which is applied only on interfaces configured as **family mpls**. All **family mpls** switch interfaces use the global EXP classifier to classify MPLS traffic.

There is no default EXP classifier. If you want to classify incoming MPLS packets using the EXP bits, you must configure a global EXP classifier. The global classifier applies to all MPLS traffic on all **family mpls** interfaces.

If a global EXP classifier is configured, MPLS traffic on **family mpls** interfaces uses the EXP classifier. If a global EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

To configure an MPLS EXP classifier using the CLI:

1. Create an EXP classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1 | exp) classifier-name forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

2. Apply the EXP classifier to the switch interfaces:

```
[edit class-of-service]
user@switch# set system-defaults classifiers exp classifier-name
```

### Related Documentation

- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 158](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\)](#)
- [Defining CoS BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\)](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 199](#)

## Configuring MPLS Firewall Filters and Policers

---

You can configure firewall filters to filter MPLS traffic. To use an MPLS firewall filter, you must first configure the filter and then apply it to an interface you have configured for forwarding MPLS traffic. 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.





**NOTE:** You can configure ingress MPLS firewall filters only. Egress MPLS firewall filters are not supported. You cannot apply MPLS firewall filters to loopback interfaces.

When you configure an MPLS firewall filter, you define filtering criteria (terms, with match conditions) for the packets and an action (action, or action modifier) for the switch to take if the packets match the filtering criteria.

- [Table 19 on page 185](#) describes the match conditions you can configure for MPLS firewall filters at the `[edit firewall family mpls filter filter-name term term-name from]` hierarchy level.



**NOTE:** If a packet has multiple MPLS labels, the filter applies the match conditions to only the bottom label in the label stack.

**Table 19: Supported Match Conditions for MPLS Firewall Filters**

| Match Condition     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>exp number</i>   | <p>Experimental (EXP) bit number or range of bit numbers in the MPLS header of a packet.</p> <p>For <i>number</i>, you can specify one or more values from 0 through 7 in binary, decimal or hexadecimal format, as given below:</p> <ul style="list-style-type: none"> <li>• A single EXP bit—for example, <b>exp 3</b></li> <li>• Several EXP bits—for example, <b>exp 0,4</b></li> <li>• A range of EXP bits—for example, <b>exp [0-5]</b></li> </ul> |
| <i>label number</i> | <p>MPLS label value or range of label values in the MPLS header of a packet.</p> <p>For <i>number</i>, you can specify one or more values from 0 through 1048575 in decimal or hexadecimal format, as given below:</p> <ul style="list-style-type: none"> <li>• A single label—for example, <b>label 3</b></li> <li>• Several labels—for example, <b>label 0,4</b></li> <li>• A range of labels—for example, <b>label [0-5]</b></li> </ul>               |

- [Table 20 on page 185](#) describes the actions you can configure for MPLS firewall filters at the `[edit firewall family mpls filter filter-name term term-name then]` hierarchy level.

**Table 20: Supported Actions for MPLS Firewall Filters**

| Action                    | Description                                                                                                                                                                                                                                                                |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>accept</i>             | Accept a packet                                                                                                                                                                                                                                                            |
| <i>count counter-name</i> | <p>Count the number of packets that pass this filter or term.</p> <p><b>NOTE:</b> We recommend that you configure a counter for each term in a firewall filter, so that you can monitor the number of packets that match the conditions specified in each filter term.</p> |

Table 20: Supported Actions for MPLS Firewall Filters (*continued*)

| Action                     | Description                                                                                                  |
|----------------------------|--------------------------------------------------------------------------------------------------------------|
| <b>discard</b>             | Discard a packet silently without sending an Internet Control Message Protocol (ICMP) message                |
| <b>policer</b>             | Starting with Junos OS 13.2X51-D15, you can send traffic matched by an MPLS filter to a two-color policer.   |
| <b>three-color-policer</b> | Starting with Junos OS 13.2X51-D15, you can send traffic matched by an MPLS filter to a three-color policer. |

- [Configuring an MPLS Firewall Filter on page 186](#)
- [Applying an MPLS Firewall Filter to an MPLS Interface on page 186](#)
- [Configuring Policers for LSPs on page 187](#)

## Configuring an MPLS Firewall Filter

To configure an MPLS firewall filter:

1. Configure the filter name, term name, and at least one match condition—for example, match on MPLS packets with EXP bits set to either 0 or 4:

```
[edit firewall family mpls]
user@switch# set filter ingress-exp-filter term term-one from exp 0,4
```

2. In each firewall filter term, specify the actions to take if the packet matches all the conditions in that term—for example, count MPLS packets with EXP bits set to either 0 or 4:

```
[edit firewall family mpls filter ingress-exp-filter term term-one then]
user@switch# set count counter0
user@switch# set accept
```

## Applying an MPLS Firewall Filter to an MPLS Interface

To apply the MPLS firewall filter to an interface you have configured for forwarding MPLS traffic (using the **family mpls** statement at the **[edit interfaces *interface-name* unit *unit-number*]** hierarchy level):



**NOTE:** You can apply firewall filters only to filter MPLS packets that enter an interface.

1. Apply the firewall filter to an MPLS interface—for example, apply the firewall filter to interface xe-0/0/5:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls filter input ingress-exp-filter
```

2. Review your configuration and issue the **commit** command:

```
[edit interfaces]
```

```
user@switch# commit
commit complete
```

## Configuring Policers for LSPs

Starting with Junos OS 13.2X51-D15, you can send traffic matched by an MPLS filter to a two-color policer or three-color policer. 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.

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.
- 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.

### Related Documentation

- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)
- [Supported MPLS Scaling Values on page 221](#)
- [Overview of Policers](#)

## Configuring MPLS to Gather Statistics

You can configure MPLS so that it periodically gathers traffic statistics about all MPLS sessions, including transit sessions, by configuring the **statistics** statement. You must configure the **statistics** statement if you want to collect MPLS traffic statistics using SNMP polling of MPLS Management Information Bases (MIBs).

To enable or disable MPLS statistics collection, include the **statistics** statement:

```
statistics {
```

```

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

```

You can configure these statements at the following hierarchy levels:

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

The default interval is 300 seconds.

If you configure the **file** option, the statistics are placed in a file, with one entry per LSP. During the specified interval, the following information is recorded in this file:

- The number of packets, number of bytes, packets per second, and bytes per second transmitted by each LSP. Feature parity for the display of packet and byte statistics for sub-LSPs of a point-to-multipoint LSP on the Junos Trio chipset is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.
- The percent of bandwidth transmitted over a given LSP in relation to the bandwidth percentage configured for that LSP. If no bandwidth is configured for an LSP, 0 percent is recorded in the percentage column.

At the end of each periodic report, a summary shows the current time, total number of sessions, number of sessions read, number of sessions ignored, and read errors, if any. Ignored sessions are typically those not in the up state or those with a reserved (0 through 15) incoming label (typically the egress point of an LSP). The reason for a read error appears on the same line as the entry for the LSP on which the error occurred. Gathering statistics is an unreliable process; occasional read errors might affect their accuracy. Sample output follows:

|                                                               |           |              |          |           |     |
|---------------------------------------------------------------|-----------|--------------|----------|-----------|-----|
| lsp6                                                          | 0 pkt     | 0 Byte       | 0 pps    | 0 Bps     | 0   |
| lsp5                                                          | 0 pkt     | 0 Byte       | 0 pps    | 0 Bps     | 0   |
| lsp6.1                                                        | 34845 pkt | 2926980 Byte | 1049 pps | 88179 Bps | 132 |
| lsp5.1                                                        | 0 pkt     | 0 Byte       | 0 pps    | 0 Bps     | 0   |
| lsp4                                                          | 0 pkt     | 0 Byte       | 0 pps    | 0 Bps     | 0   |
| Dec 7 17:28:38 Total 6 sessions: 5 success, 0 fail, 1 ignored |           |              |          |           |     |

#### Related Documentation

- [Configuring Automatic Bandwidth Allocation for LSPs on page 172](#)

## Configuring MPLS on Provider Edge Switches

To implement MPLS, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch—and at least one provider switch. You can configure

the customer edge (CE) interfaces on the PE switches of the MPLS network using IP over MPLS.

This topic describes how to configure an ingress PE switch and an egress PE switch using IP over MPLS:

1. [Configuring the Ingress PE Switch on page 189](#)
2. [Configuring the Egress PE Switch on page 190](#)

## Configuring the Ingress PE Switch

To configure the ingress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.10.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.6.1/24
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:



**NOTE:** You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
```

3. Configure OSPF traffic engineering:

```
[edit protocols ospf]
user@switch# set traffic-engineering
```

4. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols rsvp]
user@switch# set interface lo0.0
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

5. Configure MPLS traffic engineering.

```
[edit protocols mpls]
user@switch# set traffic-engineering
```

6. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

7. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
```

- ```

user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls

```
8. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:
 

```

[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 121.100.10.1/16

```
  9. Configure this Layer 3 customer edge interface for the routing protocol:
 

```

[edit]
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3.0

```
  10. Configure an LSP on the ingress PE switch (192.168.10.1) to send IP packets over MPLS to the egress PE switch (192.168.12.1):
 

```

[edit protocols mpls]
user@switch# set label-switched-path lsp_1 to 192.168.12.1

```
  11. Disable constrained-path LSP computation for this LSP:
 

```

[edit protocols mpls]
user@switch# set label-switched-path lsp_1 no-cspf

```
  12. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:
 

```

[edit routing-options]
user@switch# set static route 2.2.2.0/24 next-hop 192.168.10.1
user@switch# set static route 2.2.2.0/24 resolve

```

## Configuring the Egress PE Switch

To configure the egress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

```

[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.12.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.20.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.21.1/24

```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:



**NOTE:** You can use the switch address as an alternative to the loopback interface.

- ```

[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0

```
3. Configure RSVP on the loopback interface and the core interfaces:
 

```

[edit protocols rsvp]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface xe-0/0/5.0

```

```
user@switch# set rsvp interface xe-0/0/6.0
```

4. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
```

6. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 2.2.2.1/16
```

7. Configure this Layer 3 customer edge interface for the routing protocol:

```
[edit]
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3
```

8. Configure an LSP on the egress PE switch (192.168.12.1) to send IP packets over MPLS to the ingress PE switch (192.168.10.1):

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_2 to 192.168.10.1
```

9. Disable constrained-path LSP computation for this LSP:

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_2 no-cspf
```

10. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:

```
[edit routing-options]
user@switch# set static route 121.121.121.0/24 next-hop 192.168.12.1
user@switch# set static route 121.121.121.0/24 resolve
```

#### Related Documentation

- [MPLS Configuration Guidelines on page 220](#)
- [Configuring MPLS on Provider Switches on page 192](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 146](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 158](#)

## Configuring MPLS on Provider Switches

To implement MPLS, you must configure at least one provider switch as a transit switch for the MPLS packets.

MPLS requires the configuration of an interior gateway protocol (OSPF) and a signaling protocol (RSVP) on the core interfaces and the loopback interface of all the switches. This procedure includes the configuration of OSPF on the provider switch.

To configure the provider switch, complete the following tasks:

1. Configure OSPF on the loopback and core interfaces:



**NOTE:** You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
user@switch# set area 0.0.0.0 interface ae0
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
user@switch# set interface ae0
```

3. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols rsvp]
user@switch# set interface lo0.0
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
user@switch# set interface ae0
```

4. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 127.1.1.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switch# set ae0 unit 0 family inet address 10.1.9.2/24
```

5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
user@switch# set ae0 unit 0 family mpls
```





**NOTE:** You can configure `family mpls` on either individual interfaces or aggregated Ethernet interfaces. You cannot configure it on tagged VLAN interfaces.

#### Related Documentation

- [Configuring MPLS on Provider Edge Switches on page 188](#)
- [MPLS Configuration Guidelines on page 220](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 146](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 158](#)

## Configuring Reporting of Automatic Bandwidth Allocation Statistics for LSPs

Automatic bandwidth allocation allows an MPLS tunnel to automatically adjust its bandwidth allocation based on the volume of traffic flowing through the tunnel. You can configure the device to collect statistics related to automatic bandwidth allocation by completing the following steps:

1. To collect statistics related to automatic bandwidth allocation, configure the **auto-bandwidth** option for the **statistics** statement at the **[edit protocols mpls]** hierarchy level. These settings apply to all LSPs configured on the router on which you have also configured the **auto-bandwidth** statement at the **[edit protocols mpls label-switched-path *label-switched-path-name*]** hierarchy level.
 

```
statistics {
    auto-bandwidth;
    file filename <files number> <size size> <world-readable | no-world-readable>;
    interval seconds;
    no-transit-statistics;
    transit-statistics-polling;
}
```
2. Specify the ***filename*** for the files used to store the MPLS trace operation output using the **file** option. All files are placed in the directory **/var/log**. We recommend that you place MPLS tracing output in the file **mpls-log**.
3. Specify the maximum number of trace files using the **files *number*** option. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.
4. Specify the interval for calculating the average bandwidth usage by configuring a time in seconds using the **interval** option. You can also set the adjustment interval on a specific LSP by configuring the **interval** option at the **[edit protocols mpls label-switch-path *label-switched-path-name* statistics]** hierarchy level.



**NOTE:** To prevent unnecessary resignaling of LSPs, it is best to configure an LSP adjustment interval that is at least three times longer than the MPLS automatic bandwidth statistics interval. For example, if you configure a value of 30 seconds for the MPLS automatic bandwidth statistics interval (interval statement at the [edit protocols mpls statistics] hierarchy level), you should configure a value of at least 90 seconds for the LSP adjustment interval (adjust-interval statement at the [edit protocols mpls label-switched-path *label-switched-path-name* auto-bandwidth] hierarchy level).

- To trace automatic bandwidth allocation, include the **autobw-state** flag for the MPLS **traceoptions** statement at the [edit protocols mpls] hierarchy level.

The following configuration enables the MPLS traceoptions for automatic bandwidth allocation. The trace records are stored in a file called **auto-band-trace** (the filename is user configurable):

```
[edit protocols mpls]
traceoptions {
  file auto-band-trace size 10k files 10 world-readable;
  flag autobw-state;
}
```

- Using the **show log** command, you can display the automatic bandwidth allocation statistics file generated when you configure the *auto-bandwidth* statement. The following shows sample log file output taken from an MPLS statistics file named **auto-band-stats** on a router configured with an LSP named **E-D**. The log file shows that LSP **E-D** is operating over its reserved bandwidth limit initially. Before **Oct 30 17:14:57**, the router triggered an automatic bandwidth adjustment (you might see two sessions for an LSP undergoing an automatic bandwidth adjustment). By **Oct 30 17:16:57**, the LSP has been reestablished at a higher bandwidth and is now shown using less than 100 percent of its **Reserved Bw** (reserved bandwidth).

```
user@host> show log auto-band-stats
E-D          (LSP ID 5, Tunnel ID 6741)          209 pkt          17094 Byte          1 pps          90 Bps Util
 240.01% Reserved Bw          37 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10ct 30 17:13:57 Total 1 sessions: 1
success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)          241 pkt          19737 Byte          1 pps          88 Bps Util
 234.67% Reserved Bw          37 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10ct 30 17:14:27 Total 1 sessions: 1
success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)          276 pkt          22607 Byte          1 pps          95 Bps Util
 253.34% Reserved Bw          37 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10ct 30 17:14:57 Total 1 sessions: 1
success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)           0 pkt           0 Byte           0 pps           0 Bps Util
  0.00% Reserved Bw          37 Bps
E-D          (LSP ID 6, Tunnel ID 6741)           0 pkt           0 Byte           0 pps           0 Bps Util
  0.00% Reserved Bw          101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh 0x952c308, type 4, flags 0x0,
n_gw 1, nhid 0 to refcount 10ct 30 17:15:27 Total 2 sessions: 2 success, 0 fail, 0 ignored
E-D          (LSP ID 5, Tunnel ID 6741)           0 pkt           0 Byte           0 pps           0 Bps Util
```

```

0.00% Reserved Bw      37 Bps
E-D      (LSP ID 6, Tunnel ID 6741)      33 pkt      2695 Byte      1 pps      89 Bps Util
87.69% Reserved Bw      101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh 0x952c308, type 4, flags 0x0,
n_gw 1, nhid 0 to refcount 10ct 30 17:15:57 Total 2 sessions: 2 success, 0 fail, 0 ignored
E-D      (LSP ID 5, Tunnel ID 6741)      0 pkt      0 Byte      0 pps      0 Bps Util
0.00% Reserved Bw      37 Bps
E-D      (LSP ID 6, Tunnel ID 6741)      65 pkt      5338 Byte      1 pps      88 Bps Util
86.70% Reserved Bw      101 Bps
decr nh 0x952c224, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 1decr nh 0x952c308, type 4, flags 0x0,
n_gw 1, nhid 0 to refcount 10ct 30 17:16:27 Total 2 sessions: 2 success, 0 fail, 0 ignored
E-D      (LSP ID 6, Tunnel ID 6741)      97 pkt      7981 Byte      1 pps      88 Bps Util
86.70% Reserved Bw      101 Bps
decr nh 0x952c308, type 4, flags 0x0, n_gw 1, nhid 0 to refcount 10ct 30 17:16:57 Total 1 sessions: 1
success, 0 fail, 0 ignored

```

7. Issue the `show mpls lsp autobandwidth` command to display current information about automatic bandwidth allocation. The following shows sample output from the `show mpls lsp autobandwidth` command taken at about the same time as the log file shown previously:

```

user@host> show mpls lsp autobandwidth
Lspname      Last      Requested      Reserved      Highwater      AdjustTime LastAdjust
BW           BW           BW           mark           Left (sec)
E-D          300bps      812.005bps    812bps        1.56801kbps 294 sec      Wed Oct 30 17:15:26 2013

```

8. Issue the `file show` command to display the MPLS trace file. You need to specify the file location and file name (the file is located in `/var/log/`). The following shows sample trace file output is taken from an MPLS trace file named `auto-band-trace.0.gz` on a router configured with an LSP named `E-D`. The trace file shows that LSP `E-D` is operating over its reserved bandwidth limit initially. At `Oct 30 17:15:26`, the router triggers an automatic bandwidth adjustment (you might see two sessions for an LSP undergoing an automatic bandwidth adjustment). By `Oct 30 17:15:57`, the LSP has been reestablished at a higher bandwidth and is now shown using less than 100 percent of its **Reserved Bw** (reserved bandwidth).

```

user@host> file show /var/log/auto-band-trace.0.gz
Oct 30 17:13:57 trace_on: Tracing to "/var/log/E/auto-band-trace" started
Oct 30 17:13:57.466825 LSP E-D (id 5) new bytes arrived      2714 in 29
sec
Oct 30 17:14:27.466713 E-D      (LSP ID 5, Tunnel ID 6741)      241
pkt      19737 Byte      1 pps      88 Bps Util 234.67% Reserved Bw
      37 Bps
Oct 30 17:14:27.466962 LSP E-D (id 5, old id 5); sampled bytes      19737 >
bytes recorded      17094
Oct 30 17:14:27.467035 LSP E-D (id 5) new bytes arrived      2643 in 29
sec
Oct 30 17:14:57.466599 E-D      (LSP ID 5, Tunnel ID 6741)      276
pkt      22607 Byte      1 pps      95 Bps Util 253.34% Reserved Bw
      37 Bps
Oct 30 17:14:57.466758 LSP E-D (id 5, old id 5); sampled bytes      22607 >
bytes recorded      19737
Oct 30 17:14:57.466825 LSP E-D (id 5) new bytes arrived      2870 in 29
sec
Oct 30 17:15:26.265816 Adjust Autobw: LSP E-D (id 5) curr adj bw 300bps updated
with 812.005bps
Oct 30 17:15:26.266064 mpls LSP E-D Autobw change 512.005bps >= threshold 75bps
Oct 30 17:15:26.363372 Autobw Success: LSP E-D () (old id 5 new id 6) update
prev active bw 300 bps with 812 bps

```

```

Oct 30 17:15:26.363686 RPD_MPLS_PATH_BANDWIDTH_CHANGE: MPLS path (lsp E-D)
bandwidth changed, path bandwidth 812 bps
Oct 30 17:15:27.364751 RPD_MPLS_LSP_BANDWIDTH_CHANGE: MPLS LSP E-D bandwidth
changed, lsp bandwidth 812 bps
Oct 30 17:15:27.466849 E-D (LSP ID 5, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
37 Bps
Oct 30 17:15:27.467050 E-D (LSP ID 6, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
101 Bps
Oct 30 17:15:57.466858 E-D (LSP ID 5, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
37 Bps
Oct 30 17:15:57.467106 E-D (LSP ID 6, Tunnel ID 6741) 33
pkt 2695 Byte 1 pps 89 Bps Util 87.69% Reserved Bw
101 Bps
Oct 30 17:15:57.467201 LSP E-D (id 6, old id 5); LSP up after autobw adjustment
and active for 30 sec
Oct 30 17:15:57.467398 LSP E-D (id 6) psb bytes 2695 < bytes recorded
22607 total bytes 2695 in 30 sec
Oct 30 17:15:57.467461 First sample of the adjust interval after automatic bw
adjustment
Oct 30 17:15:57.467594 Update curr max avg bw 0bps of LSP E-D with new bw
716.225bps
Oct 30 17:16:27.466830 E-D (LSP ID 5, Tunnel ID 6741) 0
pkt 0 Byte 0 pps 0 Bps Util 0.00% Reserved Bw
37 Bps
Oct 30 17:16:27.467079 E-D (LSP ID 6, Tunnel ID 6741) 65
pkt 5338 Byte 1 pps 88 Bps Util 86.70% Reserved Bw
101 Bps
Oct 30 17:16:27.467171 LSP E-D (id 6, old id 6); sampled bytes 5338 >
bytes recorded 2695
Oct 30 17:16:27.467237 LSP E-D (id 6) new bytes arrived 2643 in 29
sec
Oct 30 17:16:57.466712 E-D (LSP ID 6, Tunnel ID 6741) 97
pkt 7981 Byte 1 pps 88 Bps Util 86.70% Reserved Bw
101 Bps
Oct 30 17:16:57.466870 LSP E-D (id 6, old id 6); sampled bytes 7981 >
bytes recorded 5338

```

- Related Documentation**
- [Configuring Automatic Bandwidth Allocation for LSPs on page 172](#)
  - [show mpls lsp autobandwidth on page 404](#)

## Configuring Static Label Switched Paths for MPLS

Configuring static label-switched paths (LSPs) for MPLS is similar to configuring static routes on individual switches. As with static routes, there is no error reporting, liveliness detection, or statistics reporting.

To configure static LSPs, configure the ingress PE switch and each provider switch along the path up to and including the egress PE switch.

For the ingress PE switch, configure which packets to tag (based on the packet's destination IP address), configure the next switch in the LSP, and the tag to apply to the packet. Manually assigned labels can have values from 0 through 1,048,575.

For the transit switches in the path, configure the next switch in the path and the tag to apply to the packet. Manually assigned labels can have values from 1,000,000 through 1,048,575.

The egress PE switch removes the label and forwards the packet to the IP destination. However, if the previous switch removed the label, the egress switch examines the packet's IP header and forwards the packet toward its IP destination.

Before you configure a static LSP, you must configure the basic components for an MPLS network:

- Configure two PE switches. See [“Configuring MPLS on Provider Edge Switches” on page 188](#).



**NOTE:** Do not configure LSPs at the [edit protocols mpls label-switched-path] hierarchy level on the PE switches.

- Configure one or more provider switches. See [“Configuring MPLS on Provider Switches” on page 192](#).

This topic describes how to configure an ingress PE switch, one or more provider switches, and an egress PE switch for static LSP:

1. [Configuring the Ingress PE Switch on page 197](#)
2. [Configuring the Provider and the Egress PE Switch on page 198](#)

## Configuring the Ingress PE Switch

To configure the ingress PE switch:

1. Configure an IP address for every core interface:

```
[edit interfaces]
user@switch# set interface-name unit logical-unit-number family inet address address
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure the name associated with the static LSP:

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name
```

3. Configure the next hop switch for the LSP:

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name ingress next-hop address-of-next-hop
```

4. Specify the address of the egress switch for the LSP:

```
[edit protocols mpls]
user@switch# set static-label-switched-path lsp-name ingress to address-of-egress-switch
```

5. Configure the new label that you want to add to the top of the label stack:

```
[edit protocols mpls]
```

```
user@switch# set static-label-switched-path lsp-name ingress push out-label
```

## Configuring the Provider and the Egress PE Switch

To configure a static LSP for MPLS on the provider and egress PE switch:

1. Configure a transit static LSP:

```
[edit protocols mpls]
```

```
user@switch# set static-label-switched-path lsp-name transit incoming-label
```

2. Configure the next hop switch for the LSP:

```
[edit protocols mpls]
```

```
user@switch# set static-label-switched-path lsp-name transit incoming-label next-hop  
address-of-next-hop
```

3. Only for provider switches, remove the label at the top of the label stack and replace it with the specified label:

```
[edit protocols mpls]
```

```
user@switch# set static-label-switched-path lsp-name transit incoming-label swap out-label
```

4. Only for the egress PE switch, remove the label at the top of the label stack:



**NOTE:** 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).

```
[edit protocols mpls]
```

```
user@switch# set static-label-switched-path lsp-name transit incoming-label pop
```

### Related Documentation

- [Configuring MPLS on Provider Edge Switches on page 188](#)
- [Configuring MPLS on Provider Switches on page 192](#)
- [Understanding MPLS Label Operations on page 150](#)

## Configuring Rewrite Rules for MPLS EXP Classifiers

You configure EXP rewrite rules to alter CoS values in outgoing MPLS packets on the outbound **family mpls** interfaces of a switch to match the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

To configure an EXP CoS rewrite rule, create the rule by giving it a name and associating it with a forwarding class, loss priority, and code point. This creates a rewrite table. After the rewrite rule is created, enable it on a logical **family mpls** interface. EXP rewrite rules can only be enabled on logical **family mpls** interfaces, not on physical interfaces or on interfaces of other family types. You can also apply an existing EXP rewrite rule on a logical interface.



**NOTE:** There are no default rewrite rules.

You can configure up to 64 EXP rewrite rules, but you can only use 16 EXP rewrite rules at any time on the switch. On a given **family mpls** logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.



**NOTE:** To replace an existing rewrite rule on the interface with a new rewrite rule of the same type, first explicitly remove the existing rewrite rule and then apply the new rule.

To create an EXP rewrite rule for MPLS traffic and enable it on a logical interface:

1. Create an EXP rewrite rule:

```
user@switch# set class-of-service rewrite-rules exp rewrite-rule-name forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

For example, to configure an EXP rewrite rule named **exp-rr-1** for a forwarding class named **mpls-1** with a loss priority of **low** that rewrites the EXP code point value to **001**:

```
user@switch# set class-of-service rewrite-rules exp exp-rr-1 forwarding-class mpls-1
loss-priority low code-points 001
```

2. Apply the rewrite rule to a logical interface:

```
user@switch # set class-of-service interfaces interface-name unit logical-unit rewrite-rules
exp rewrite-rule-name
```

For example, to apply a rewrite rule named **exp-rr-1** to logical interface **xe-0/0/10.0**:

```
user@switch# set class-of-service interfaces xe-0/0/10 unit 0 rewrite-rules exp exp-rr-1
```

---



**NOTE:** In this example, all forwarding classes assigned to port **xe-0/0/10** must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same interface.

---

**Related  
Documentation**

- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 158](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces](#)
- [Monitoring CoS Rewrite Rules](#)
- [Defining CoS Rewrite Rules](#)

---

## Example: Configuring MPLS-Based Layer 3 VPNs

---

You can implement an MPLS-based Layer 3 virtual private network (VPN) on QFX switches to interconnect sites for customers who want the service provider to handle all the Layer 3 routing functions. To support an MPLS-based Layer 3 VPN, you need to add components of the Layer 3 VPN to the configuration of the two provider edge (PE) switches. You do not need to change the configuration of the provider switches.

This example shows how to configure an MPLS-based Layer 3 VPN spanning two corporate sites:

- [Requirements on page 201](#)
- [Overview and Topology on page 201](#)
- [Configuring the Local PE Switch on page 204](#)
- [Configuring the Remote PE Switch on page 206](#)



## Requirements

This example uses the following software and hardware components:

- Junos OS Release 12.3 or later for the QFX Series
- Three QFX switches

Before you configure the Layer 3 VPN components, you must configure the basic components for an MPLS network:

- Configure two PE switches. See [“Configuring MPLS on Provider Edge Switches” on page 188](#).
- Configure one or more provider switches. See [“Configuring MPLS on Provider Switches” on page 192](#).

## Overview and Topology

Layer 3 VPNs allow customers to leverage the service provider’s technical expertise to ensure efficient site-to-site routing. The customer’s customer edge (CE) switch uses a routing protocol such as BGP or OSPF to communicate with the service provider’s provider edge (PE) switch to carry IP prefixes across the network. MPLS-based Layer 3 VPNs use only IP over MPLS; other protocol packets are not supported. This example includes two PE switches, PE1 and PE2.


In the basic MPLS configuration of the PE switches using IP over MPLS, the PE switches were configured to use OSPF as the routing protocol between the MPLS switches and RSVP as the signaling protocol. Traffic engineering was enabled. A label-switched path (LSP) was configured.

The following components must be added to the PE switches for an MPLS-based Layer 3 VPN:

- BGP group with **family inet-vpn unicast**
- Routing instance with instance type **vrf**

[Figure 11 on page 201](#) shows the topology used in this example.

**Figure 11: Configuring an MPLS-Based Layer 3 VPN**



[Table 21 on page 202](#) shows the settings of the customer edge interface on the local CE switch.

Table 21: Local CE Switch in the MPLS-Based Layer 3 VPN Topology

| Property                 | Settings                                                                                                            | Description                         |
|--------------------------|---------------------------------------------------------------------------------------------------------------------|-------------------------------------|
| Local CE switch hardware | QFX switch                                                                                                          | CE1                                 |
| Customer edge interface  | <b>ge-0/0/14 unit 0</b><br><b>family inet</b><br><b>address 51.51.0.14/16</b><br>protocols ospf interface ge-0/0/14 | Interface that connects CE1 to PE1. |

Table 22 on page 202 shows the settings of the customer edge interface on the remote CE switch.

Table 22: Remote CE Switch in the MPLS-Based Layer 3 VPN Topology

| Property                  | Settings                                                                                                            | Description                         |
|---------------------------|---------------------------------------------------------------------------------------------------------------------|-------------------------------------|
| Remote CE switch hardware | QFX switch                                                                                                          | CE2                                 |
| Customer edge interface   | <b>ge-0/0/14 unit 0</b><br><b>family inet</b><br><b>address 11.22.26.1/16</b><br>protocols ospf interface ge-0/0/14 | Interface that connects CE2 to PE2. |

Table 23 on page 202 shows the Layer 3 VPN components of the local PE switch.

Table 23: Layer 3 VPN Components of the Local PE Switch

| Property                 | Settings                                                                                | Description                                                                                                                                                                                                                                                                       |
|--------------------------|-----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local PE switch hardware | QFX switch                                                                              | PE1                                                                                                                                                                                                                                                                               |
| Customer edge interface  | <b>ge-5/0/24 unit 0</b><br><b>family inet</b><br><b>address 51.51.0.1/16</b>            | Connects PE1 to CE1.<br><br><b>NOTE:</b> The <b>family inet</b> configuration should already have been completed as part of the basic MPLS configuration of the PE switch for IP over MPLS. It is included here to show what was specified for that portion of the configuration. |
| Core interface           | <b>xe-0/0/6 unit 0</b><br><b>family inet address 60.0.0.60/16</b><br><b>family mpls</b> | Connects PE1 to P.<br><br><b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.                                        |

Table 23: Layer 3 VPN Components of the Local PE Switch (*continued*)

| Property           | Settings                                         | Description                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Loopback interface | lo0 unit 0<br>family inet address 21.21.21.21/32 | <b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration. |
| BGP                | bgp                                              | Added for the Layer 3 VPN configuration.                                                                                                                                                                         |
| Routing instance   | L3VPN-1                                          | Added for the Layer 3 VPN configuration.                                                                                                                                                                         |

Table 24 on page 203 shows the Layer 3 VPN components of the remote PE switch.

Table 24: Layer 3 VPN Components of the Remote PE Switch

| Property                  | Settings                                                           | Description                                                                                                                                                                                                                                                                                                                                            |
|---------------------------|--------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Remote PE switch hardware | QFX switch                                                         | PE2                                                                                                                                                                                                                                                                                                                                                    |
| Customer edge interface   | ge-0/0/14 unit 0<br>family inet<br>address 11.22.26.14/16          | Connects PE2 to CE2.<br><br>For the Layer 3 VPN configuration, added <b>family mpls</b> .<br><br><b>NOTE:</b> The <b>family inet</b> configuration should already have been completed as part of the basic MPLS configuration of the PE switch for IP over MPLS. It is included here to show what was specified for that portion of the configuration. |
| Core interface            | xe-0/0/6 unit 0<br>family inet address 60.2.0.60/16<br>family mpls | Connects PE1 to P.<br><br><b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.                                                                                                             |
| Loopback interface        | lo0 unit 0<br>family inet address 22.22.22.22/32                   | <b>NOTE:</b> This portion of the configuration should already have been completed as part of the basic MPLS configuration. It is included here to show what was specified for that portion of the configuration.                                                                                                                                       |
| BGP                       | bgp                                                                | Added for the Layer 3 VPN configuration.                                                                                                                                                                                                                                                                                                               |
| Routing instances         | L3VPN-1                                                            | Added for the Layer 3 VPN configuration.                                                                                                                                                                                                                                                                                                               |

## Configuring the Local PE Switch

**CLI Quick Configuration** To quickly configure the Layer 3 VPN components on the local PE switch, copy the following commands and paste them into the switch terminal window of PE1:

```
[edit]
set protocols bgp local-address 21.21.21.21 family inet-vpn unicast
set protocols bgp group PE1-PE2 type internal
set protocols bgp neighbor 22.22.22.22
set routing-instances L3VPN-1 instance-type vrf
set routing-instances L3VPN-1 description "BETWEEN PE1 AND PE2"
set routing-instances L3VPN-1 interface ge-5/0/24
set routing-instances L3VPN-1 protocols ospf interface ge-5/0/24
set routing-instances L3VPN-1 route-distinguisher 21:21
set routing-instances L3VPN-1 vrf-target target:21:21
set routing-instances L3VPN-1 vrf-table-label
set routing-options router-id 21.21.21.21
set routing-options autonomous-system 10
```

**Step-by-Step Procedure** To configure the Layer 3 VPN components on the local PE switch:

1. Configure BGP, specifying the loopback address as the local address and specifying **family inet-vpn unicast**:

```
[edit protocols bgp]
user@switchPE1# set local-address 21.21.21.21 family inet-vpn unicast
```

2. Configure the BGP group, specifying the group name and type:

```
[edit protocols bgp]
user@switchPE1# set group PE1-PE2 type internal
```

3. Configure the BGP neighbor, specifying the loopback address of the remote PE switch as the neighbor's address:

```
[edit protocols bgp]
user@switchPE1# set neighbor 22.22.22.22
```

4. Configure the routing instance, specifying the routing-instance name and using **vrf** as the instance type:

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 instance-type vrf
```

5. Configure a description for this routing instance:

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 description "BETWEEN PE1 AND PE2"
```

6. Configure the routing instance for the OSPF interface:

```
[edit routing-instances]
user@switchPE2# set L3VPN-1 protocols ospf interface ge-5/0/24
```

7. Configure the routing instance to use a route distinguisher:

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 route-distinguisher 21:21
```



**NOTE:** Each routing instance that you configure on a PE switch must have a unique route distinguisher associated with it. VPN routing instances require a route distinguisher to allow BGP to distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPNs. If you configure different VPN routing instances with the same route distinguisher, the commit fails.

8. Configure the VPN routing and forwarding (VRF) target of the routing instance:

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 vrf-target target:21:21
```



**NOTE:** You can create more complex policies by explicitly configuring VRF import and export policies using the import and export options. See the *Junos OS VPNs Library for Routing Devices*.

9. Configure this routing instance with **vrf-table-label**, which maps the inner label of a packet to a specific VPN routing and forwarding (VRF) table and allows the examination of the encapsulated IP header:

```
[edit routing-instances]
user@switchPE1# set L3VPN-1 vrf-table-label
```

10. Configure the router ID and autonomous system (AS):



**NOTE:** We recommend that you explicitly configure the router identifier under the [edit routing-options] hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.

```
[edit routing-options]
user@switchPE1# set router-id 21.21.21.21 autonomous-system 10
```

**Results** Display the results of the configuration:

```
user@switchPE1> show configuration
```

```
interfaces {
  ge-0/0/14 {
    unit 0 {
      family inet {
        address 51.51.0.1/16;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 21.21.21.21/32;
      }
    }
  }
}
```

```

    }
  }
}
xe-0/0/6 {
  unit 0 {
    family inet {
      address 60.0.0.60/16;
    }
    family mpls;
  }
}
protocols {
  mpls {
    label-switched-path 21-22 {
      from 21.21.21.21;
      to 22.22.22.22;
      no-cspf;
    }
    interface xe-0/0/6.0;
    interface lo0.0;
  }
  bgp {
    local-address 21.21.21.21;
    family inet-vpn {
      unicast;
    }
    group PE1-PE2 {
      type internal;
      neighbor 22.22.22.22;
    }
  }
  ospf
    traffic-engineering;
    area 0.0.0.0 {
      interface lo0.0;
      interface xe-0/0/6.0;
    }
  }
}
routing-instances {
  L3VPN-1 {
    instance-type vrf;
    description "BETWEEN PE1 AND PE2";
    route-distinguisher 21:21;
    vrf-target target:21:21;
    vrf-table-label;
  }
  routing-options {
    router-id 21.21.21.21;
    autonomous-system 10;
  }
}

```

## Configuring the Remote PE Switch

**CLI Quick Configuration** To quickly configure the Layer 3 VPN components on the remote PE switch, copy the following commands and paste them into the switch terminal window of PE2:

[edit]

```

set protocols bgp local-address 22.22.22.22 family inet-vpn unicast
set protocols bgp group PE1-PE2 type internal
set protocols bgp neighbor 21.21.21.21
set routing-instances L3VPN-1 instance-type vrf
set routing-instances L3VPN-1 description "BETWEEN PE1 AND PE2"
set routing-instances L3VPN-1 interface ge-11/0/14.0
set routing-instances L3VPN-1 protocols ospf interface ge-11/0/14.0
set routing-instances L3VPN-1 route-distinguisher 21:21
set routing-instances L3VPN-1 vrf-target target:21:21
set routing-instances L3VPN-1 vrf-table-label;
set routing-options router-id 22.22.22.22
set routing-options autonomous-system 10

```

### Step-by-Step Procedure

To configure Layer 3 VPN components on the remote PE switch:

1. Configure BGP, specifying the loopback address as the local address and specifying **family inet-vpn unicast**:  

```

[edit protocols bgp]
user@switchPE2# set local-address 22.22.22.22 family inet-vpn unicast

```
2. Configure the BGP group, specifying the group name and type:  

```

[edit protocols bgp]
user@switchPE2# set group PE1-PE2 type internal

```
3. Configure the BGP neighbor, specifying the loopback address of the remote PE switch as the neighbor's address:  

```

[edit protocols bgp]
user@switchPE2# set neighbor 21.21.21.21

```
4. Configure the routing instance, specifying the routing-instance name and using **vrf** as the instance type:  

```

[edit routing-instances]
user@switchPE2# set L3VPN-1 instance-type vrf

```
5. Configure a description for this routing instance:  

```

[edit routing-instances]
user@switchPE1# set L3VPN-1 description "BETWEEN PE1 AND PE2"

```
6. Configure the routing instance to apply to the customer edge interface:  

```

[edit routing-instances]
user@switchPE2# set L3VPN-1 interface ge-0/0/14.0

```
7. Configure the routing instance for the OSPF interface:  

```

[edit routing-instances]
user@switchPE2# set L3VPN-1 protocols ospf interface ge-11/0/14.0

```
8. Configure the routing instance to use a route distinguisher, using the format *ip-address:number*:  

```

[edit routing-instances]
user@switchPE2# set L3VPN-1 route-distinguisher 21:21

```
9. Configure the VPN routing and forwarding (VRF) target of the routing instance:  

```

[edit routing-instances]
user@switchPE2# set L3VPN-1 vrf-target target:21:21

```
10. Configure this routing instance with **vrf-table-label**, which maps the inner label of a packet to a specific VPN routing and forwarding (VRF) table and allows the examination of the encapsulated IP header.  

```

[edit routing-instances]

```

- ```

user@switchPE2# set L3VPN-1 vrf-label-label
11. Configure the router ID and autonomous system (AS):
[edit routing-options]
user@switchPE2# set router-id 22.22.22.22 autonomous-system 10

```

**Results** Display the results of the configuration:

```

user@switchPE2> show configuration
interfaces {
  ge-0/0/14 {
    unit 0 {
      family inet {
        address 11.22.26.14/16;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 22.22.22.22/32;
      }
    }
  }
  xe-0/0/6 {
    unit 0 {
      family inet {
        address 60.2.0.60/16;
      }
      family mpls;
    }
  }
}
protocols {
  mpls {
    label-switched-path 22-21 {
      from 22.22.22.22;
      to 21.21.21.21;
      no-cspf;
    }
    interface xe-0/0/6.0;
    interface lo0.0;
  }
  bgp {
    local-address 22.22.22.22;
    family inet-vpn {
      unicast;
    }
    group PE1-PE2 {
      type internal;
      neighbor 21.21.21.21;
    }
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface ge-0/0/14.0;
      interface lo0.0;
    }
  }
}

```



```

        interface xe-0/0/6.0;
    }
}
routing-instances {
    L3VPN-1 {
        instance-type vrf;
        description "BETWEEN PE1 AND PE2";
        route-distinguisher 21:21;
        vrf-target target:21:21;
        vrf-table-label;
    }
    routing-options {
        router-id 22.22.22.22;
        autonomous-system 10;
    }
}

```

- Related Documentation**
- [Configuring MPLS on Provider Edge Switches on page 188](#)
  - [Configuring MPLS on Provider Switches on page 192](#)

## Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks

This example shows how to configure 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 209](#)
- [Overview on page 209](#)
- [Configuration on page 212](#)
- [Verification on page 217](#)

### 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 12 on page 210](#), PE1 and PE2 are dual-stack BGP routers or switches, meaning they have both IPv4 and IPv6 stacks. The PE devices link the IPv6 networks through the customer edge (CE) routers or switches to the IPv4 core network. The CE devices and the PE devices connect through a link layer that can carry IPv6 traffic. The PE devices 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 12: IPv6 Networks Linked by MPLS IPv4 Tunnels**



The two PE devices 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 devices can learn the IPv6 routes from the CE devices connected to them using 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 device and CE device 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 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 devices 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 device is not a Juniper Networks routing or switching 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 devices 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 device in [Figure 12 on page 210](#) receives an IPv6 packet from the CE1 device, it performs a lookup in the IPv6 forwarding table. If the destination matches a prefix learned from the CE2 device, then no labels need to be pushed and the packet is simply sent to the CE2 device. If the destination matches a prefix that was learned from the PE2 device, 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 xe-0/0/5 unit 2 family inet6 address ::10.1.1.2/126
set interfaces xe-0/0/5 unit 2 family mpls
set interfaces xe-0/0/6 unit 5 family inet address 10.1.1.5/30
set interfaces xe-0/0/6 unit 5 family inet6
set interfaces xe-0/0/6 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 xe-0/0/5.2
set protocols mpls interface xe-0/0/6.5
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 xe-0/0/6.5
set protocols ospf area 0.0.0.0 interface lo0.2 passive
set protocols rsvp interface xe-0/0/6.5
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 xe-0/0/5 unit 10 family inet address 10.1.1.10/30
set interfaces xe-0/0/5 unit 10 family inet6
set interfaces xe-0/0/5 unit 10 family mpls
set interfaces xe-0/0/6 unit 13 family inet6 address ::10.1.1.13/126
set interfaces xe-0/0/6 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 xe-0/0/5.10
set protocols mpls interface xe-0/0/6.13
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 xe-0/0/5.10
set protocols ospf area 0.0.0.0 interface lo0.4 passive
set protocols rsvp interface xe-0/0/5.10
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 xe-0/0/5 unit 6 family inet address 10.1.1.6/30
set interfaces xe-0/0/5 unit 6 family inet6
set interfaces xe-0/0/5 unit 6 family mpls
set interfaces xe-0/0/6 unit 9 family inet address 10.1.1.9/30
set interfaces xe-0/0/6 unit 9 family inet6
set interfaces xe-0/0/6 unit 9 family mpls
set interfaces lo0 unit 3 family inet address 1.1.1.3/32
set protocols mpls interface xe-0/0/5.6
set protocols mpls interface xe-0/0/6.9
set protocols ospf area 0.0.0.0 interface xe-0/0/5.6
set protocols ospf area 0.0.0.0 interface xe-0/0/6.9
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set protocols rsvp interface xe-0/0/5.6
set protocols rsvp interface xe-0/0/6.9
set routing-options router-id 1.1.1.3
set routing-options autonomous-system 2
```

```

Device CE1    set interfaces xe-0/0/5 unit 1 family inet6 address ::10.1.1.1/126
               set interfaces xe-0/0/5 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 xe-0/0/5 unit 14 family inet6 address ::10.1.1.14/126
               set interfaces xe-0/0/5 unit 14 family mpls
               set interfaces lo0 unit 5 family inet6 address ::1.1.1.5/128
               set protocols bgp group toPE2 type external
               set protocols bgp group toPE2 local-address ::10.1.1.14
               set protocols bgp group toPE2 family inet6 unicast
               set protocols bgp group toPE2 export send-v6
               set protocols bgp group toPE2 peer-as 2
               set protocols bgp group toPE2 neighbor ::10.1.1.13
               set policy-options policy-statement send-v6 from family inet6
               set policy-options policy-statement send-v6 from protocol direct
               set policy-options policy-statement send-v6 then accept
               set routing-options router-id 1.1.1.5
               set routing-options autonomous-system 3

```

### Configuring Device PE1

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

To configure Device PE1:

1. Configure the interfaces.
 

```

[edit interfaces]
user@PE1# set xe-0/0/5 unit 2 family inet6 address ::10.1.1.2/126
user@PE1# set xe-0/0/5 unit 2 family mpls

user@PE1# set xe-0/0/6 unit 5 family inet address 10.1.1.5/30
user@PE1# set xe-0/0/6 unit 5 family inet6
user@PE1# set xe-0/0/6 unit 5 family mpls

user@PE1# set lo0 unit 2 family inet address 1.1.1.2/32

```
2. Configure MPLS on the interfaces.
 

```

[edit protocols mpls]
user@PE1# set ipv6-tunneling
user@PE1# set interface xe-0/0/5.2

```

```
user@PE1# set interface xe-0/0/6.5
```

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 xe-0/0/6.5
user@PE1# set interface lo0.2 passive
```

5. Configure a signaling protocol.

```
[edit protocols]
user@PE1# set rsvp interface xe-0/0/6.5
```

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
xe-0/0/5 {
  unit 2 {
```

```
        family inet6 {
            address ::10.1.1.2/126;
        }
        family mpls;
    }
}
xe-0/0/6 {
    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 xe-0/0/5.2;
    interface xe-0/0/6.5;
}
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 xe-0/0/6.5;
interface lo0.2 {
passive;
}
}
}
}
rsvp {
interface xe-0/0/6.5;
}
}

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 212](#).

## 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**

## Verifying That MPLS Is Working Correctly

To verify that MPLS is working correctly, perform the following tasks:

1. [Verifying the Physical Layer on the Switches on page 218](#)
2. [Verifying the Routing Protocol on page 218](#)
3. [Verifying the Core Interfaces Being Used for the MPLS Traffic on page 219](#)
4. [Verifying RSVP on page 219](#)

## Verifying the Physical Layer on the Switches

**Purpose** Verify that the interfaces are up. Perform this verification task on each of the switches.

**Action** user@switch> **show interfaces xe-\* terse**

Interface	Admin	Link	Proto	Local	Remote
xe-0/0/0	up	up			
xe-0/0/0.0	up	up			
xe-0/0/1.0	up	up			
xe-0/0/2.0	up	up			
xe-0/0/3.0	up	up	inet	2.2.2.1/16	
xe-0/0/4.0	up	up			
xe-0/0/5.0	up	up	inet mpls	10.1.5.1/24	
xe-0/0/6.0	up	up	inet mpls	10.1.6.1/24	

**Meaning** The **show interfaces terse** command displays status information about the 10-Gigabit Ethernet interfaces on the switch. This output verifies that the interfaces are **up**. The output for the protocol family (Proto column) of the core interfaces (xe-0/0/5.0 and xe-0/0/6.0), shows that these interfaces are configured as both **inet** and **mpls**. The **Local** column for the core interfaces shows the IP address configured for these interfaces.

## Verifying the Routing Protocol

**Purpose** Verify the state of the configured routing protocol. You should perform this verification task on each of the switches. The state should be **Full**. If you have configured OSPF as the routing protocol, use the **show ospf neighbor** command to verify that the routing protocol is communicating with the switch neighbors.

**Action** user@switch> **show ospf neighbor**

Address	Interface	State	ID	Pri	Dead
127.1.1.1	xe-0/0/5	Full	10.10.10.10	128	39

**Meaning** The **show ospf neighbor** command displays the status of the routing protocol that has been configured on this switch. The output shows that the state is **Full**, meaning that the routing protocol is operating correctly—that is, hello packets are being exchanged between directly connected neighbors. For additional information on checking and monitoring routing protocols, see the [Junos OS Routing Protocols and Policies Command Reference](#).

## Verifying the Core Interfaces Being Used for the MPLS Traffic

**Purpose** Verify that the state of the MPLS interface is **Up**. You should perform this verification task on each of the switches.

**Action** user@switch> **show mpls interface**

Interface	State	Administrative groups
ge-0/0/5	Up	<none>
ge-0/0/6	Up	<none>

**Meaning** The **show mpls interface** command displays the status of the core interfaces that have been configured to belong to **family mpls**. This output shows that the interface configured to belong to **family mpls** is up.

## Verifying RSVP

**Purpose** Verify the state of the RSVP session. You should perform this verification task on each of the switches.

user@switch> **show rsvp session**

Ingress RSVP: 1 sessions

To	From	State	Rt	Style	Labelin	Labelout	LSPname
127.1.1.3	127.1.1.1	Up	0	1 FF	-	300064	lsp_to_pe2_ge1

Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions

To	From	State	Rt	Style	Labelin	Labelout	LSPname
127.1.1.1	127.1.1.3	Up	0	1 FF	299968	-	lsp_to_pe1_ge1

Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions  
Total 0 displayed, Up 0, Down 0

**Meaning** This output confirms that the RSVP sessions are up.

**Related Documentation**

- [Configuring MPLS on Provider Edge Switches on page 188](#)

- [Configuring MPLS on Provider Switches on page 192](#)

## MPLS Configuration Guidelines

---

When configuring MPLS on QFX Series devices or on EX4600, note that the number of IP prefixes supported depends on the specific platform being used. See the scale specifications in the data sheet of your device for additional information.

- We recommend the following:
  - If your ingress provider edge (PE) switch needs to support more than 8000 external IP prefixes, use a larger capacity device as an ingress PE switch.
  - If you use a switch as a route reflector for BGP labeled routes, use it as a dedicated route reflector (that is, the switch must not participate in managing data traffic).
  - If you use a switch as a PE switch or as a route reflector for BGP labeled routes, configure routing policies on the PE switch and the route reflector to filter external IP routes from the routing table.

The configuration example for a routing policy named `fib_policy` (at the `[edit policy-options` and `[edit routing-options` hierarchy levels) to filter BGP labeled routes from the `inet.0` routing table is given below:

```
user@switch# show policy-options
policy-statement fib_policy {
  from {
    protocol bgp;
    rib inet.0;
  }
  then reject;
}

user@switch# show routing-options
forwarding-table {
  export fib_policy;
}
```

- Packet fragmentation using the `allow-fragmentation` statement at the `[edit protocols mpls path-mtu]` hierarchy level is not supported on QFX Series devices or on the EX4600 switch. Therefore, you must ensure that the maximum transmission unit (MTU) values configured on every MPLS interface is sufficient to handle MPLS packets. The packets whose size exceeds the MTU value of an interface will be dropped.

### Related Documentation

- [Configuring MPLS on Provider Edge Switches on page 188](#)
- [Configuring MPLS on Provider Switches on page 192](#)
- [Configuring a Global MPLS EXP Classifier on page 184](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 199](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)

## Supported MPLS Scaling Values

This topic lists the MPLS scaling values supported on QFX Series switches.

[Table 25 on page 221](#) lists the MPLS scaling values supported on Juniper QFX switches and on the EX4600 switch.

**Table 25: MPLS Scaling Values**

Feature	QFX3500 Scaling Value	QFX5100 and EX4600 Scaling Value	QFX10002 Scaling Value
Maximum number of MPLS labels in a packet's label stack	3 labels for Push operations	3 labels for Push operations	5 labels for Push operations
	2 labels for Pop operations	2 labels for Pop operations	8 labels for Pop operations
	1 label for Swap operations	1 label for Swap operations	1 label for Swap operations
Maximum number of MPLS labels on provider switches	4096	16386	80000 (Junos limit)
Maximum number of tunnel (combination of routes and LSPs) initiations	Ingress LSPs: 1024	Ingress LSPs: 1024	Ingress LSPs: 32000 (Junos limit)
	Transit LSPs: 4000	Transit LSPs: 16386	Transit LSPs: 80000 (Junos limit)
Maximum number of unique next-hops on egress provider edge (PE) switches	512	512	
Maximum number of MPLS firewall filters	768	1536	8000 ingress 8000 egress
Virtual Routing and Forwarding (VRF)	1K	1K	4K
Layer 3 Host	IPv4: 8K	See <i>Understanding the Unified Forwarding Table</i> .	
Layer 3 Longest Prefix Match (LPM)	IPv4: 16K	See <i>Understanding the Unified Forwarding Table</i> .	
	IPv6: 4K		

- Related Documentation**
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)
  - [MPLS Configuration Guidelines on page 220](#)

## MPLS Stitching For Virtual Machine Connection

By using MPLS, the stitching feature of Junos OS provides connectivity between virtual machines that reside either on opposite sides of data center routers or in different data centers. An external controller, programmed in the data-plane, assigns MPLS labels to both virtual machines and servers. Then, the signaled MPLS labels are used between the data center routers, generating static link switched paths (LSPs), resolved over BGP labeled unicast, RSVP or LDP, to provide the routes dictated by the labels.


- [When Would I Use Stitching? on page 222](#)
- [How Does MPLS Stitching Work? on page 222](#)
- [How Do I Configure Stitching? on page 223](#)
- [Which Switches Support Stitching? on page 223](#)
- [Q&A on page 223](#)

### When Would I Use Stitching?

There are several ways to connect virtual machines. One option when you have virtual machines on opposite sides of a router (or different data centers) is to use MPLS stitching. A typical topology for using MPLS stitching is shown in [Figure 13 on page 222](#).

**Figure 13: Virtual Machines on Either Side of Routers**

ERROR: Unresolved graphic fileref="" not found in  
 "/cmsxml/default/main/supplemental/STAGING/images/".



The above topology consists of the following MPLS layers: VMs | Servers | ToRs | Router  
 ..... Router | ToRs | Servers | VMs



**NOTE:** The label on the left is the top of the label stack.

### How Does MPLS Stitching Work?

With stitching, the MPLS static allocation of labels demultiplexes incoming traffic onto any device/entity in the next layer in the direction of traffic flow. Essentially, there is a label hierarchy that picks up labels for the correct top-of-rack switch, server, and virtual machine that receives traffic. Static label assignments are done between the top-of-rack switches and the virtual machines.

For example, imagine that traffic is sent from VM1 to VM3 in [Figure 13 on page 222](#). When traffic exits Server1, its label stack is L1 | L2 | L3 where:

- L1 represents the egress top-of-rack switch ToR1.

- L2 represents the physical server, Server2, towards which the egress-side ToR will forward the traffic.
- L3: represents the virtual machine on Server2 to which the Server2 should deliver the traffic.

When traffic arrives at ToR1, it needs to be sent to ToR2. Since ToR1 and ToR2 are not directly connected, traffic must flow from ToR1 to ToR2 using label-switching starting on the outermost (top) label. Stitching has been added to static-LSP functionality to SWAP L1 to a L-BGP label that ToR2 advertises to ToR1. The label stack now must contain another label at the top to enable forwarding of the labeled packets between ToR1 and ToR2. An L-Top label is added if L-BGP is resolved over RSVP/LDP. If static LSP is resolved over L-BGP, then the top label is swapped with the L-BGP label and there is no L-Top label. When the traffic exits ToR1, the stack is: L-top | L-BGP | L2 | L3.

Traffic from ToR1 to ToR2 is then label switched over any signaled LSP.

When traffic arrives at ToR2, the top label is removed with PHP (popped) and the label stack becomes L-BGP | L2 | L3. Since L-BGP is a implicit null label, ToR2 pops the static LSP label L2 that corresponds to the egress server and then forward the packet to the egress server using the static-LSP configuration on ToR2, which corresponds to a single-hop implicit-NULL LSP.

The outgoing stack becomes L3 and the next-hop is the egress server Server2.

When traffic arrives at the egress server Server2, Server2 pops L3 and delivers the packet to VM3.

## How Do I Configure Stitching?

The new keyword **stitch** for LSPs under the command *transit* has been added to resolve the remote next-hop. The **show mpls static-lsp** command has been extended to show the LSP state as 'InProgress' whenever the LSP is waiting for protocol next-hop resolution by resolver.

## Which Switches Support Stitching?

QFX5100, QFX3500 and EX4600 support the static LSP stitching feature.

## Q&A

Q: Is link and node protection for the next-hop provided by MPLS stitching?

A: link and node protection for the next-hop of transit LSP stitched to L-BGP LSP are not needed. That is provided by L-BGP LSP.

Q: Does stitching work with real L-BGP labels?

A: No, stitching does not work when the L-BGP label is a real label.

## Related Documentation

- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)





## CHAPTER 5

# Configuration Statements for MPLS

- [adaptive](#) on page 228
- [adjust-interval](#) on page 228
- [adjust-threshold](#) on page 229
- [adjust-threshold-overflow-limit](#) on page 229
- [adjust-threshold-underflow-limit](#) on page 230
- [admin-down](#) on page 230
- [advertisement-hold-time](#) on page 231
- [associate-backup-pe-groups](#) on page 231
- [auto-bandwidth \(MPLS Tunnel\)](#) on page 232
- [backup-pe-group](#) on page 233
- [bandwidth \(Fast Reroute, Signaled, and Multiclass LSPs\)](#) on page 234
- [bandwidth-model](#) on page 235
- [bypass \(Static LSP\)](#) on page 236
- [chained-composite-next-hop](#) on page 237
- [class-of-service \(Protocols MPLS\)](#) on page 239
- [corouted-bidirectional](#) on page 240
- [corouted-bidirectional-passive](#) on page 240
- [description \(Protocols MPLS\)](#) on page 241
- [diffserv-te](#) on page 242
- [disable \(Protocols MPLS\)](#) on page 243
- [exclude \(for Fast Reroute\)](#) on page 244
- [exclude-srlg](#) on page 245
- [explicit-null \(Protocols MPLS\)](#) on page 246
- [fast-reroute \(Protocols MPLS\)](#) on page 247
- [forwarding-table](#) on page 248
- [from \(Protocols MPLS\)](#) on page 248
- [gpipid](#) on page 249
- [hop-limit](#) on page 250

- [include-all \(for Fast Reroute\) on page 251](#)
- [include-any \(for Fast Reroute\) on page 251](#)
- [ingress \(LSP\) on page 252](#)
- [install \(Protocols MPLS\) on page 253](#)
- [interface \(Protocols MPLS\) on page 254](#)
- [ipv6-tunneling on page 255](#)
- [l2circuit on page 256](#)
- [label-switched-path \(Protocols MPLS\) on page 258](#)
- [ldp-tunneling on page 260](#)
- [link-protection \(Static LSPs\) on page 261](#)
- [log-updown \(Protocols MPLS\) on page 262](#)
- [lsp-attributes on page 263](#)
- [maximum-bandwidth \(Protocols MPLS\) on page 264](#)
- [metric \(Protocols MPLS\) on page 264](#)
- [minimum-bandwidth on page 265](#)
- [monitor-bandwidth on page 265](#)
- [mtu-signaling on page 266](#)
- [no-cspf on page 267](#)
- [no-decrement-ttl on page 268](#)
- [no-install-to-address on page 269](#)
- [no-propagate-ttl on page 270](#)
- [record on page 271](#)
- [no-trap on page 272](#)
- [node-link-protection \(Protocols MPLS\) on page 273](#)
- [oam \(Protocols MPLS\) on page 274](#)
- [optimize-aggressive on page 275](#)
- [optimize-hold-dead-delay on page 276](#)
- [optimize-switchover-delay on page 277](#)
- [optimize-timer \(Protocols MPLS\) on page 278](#)
- [p2mp \(Protocols MPLS\) on page 279](#)
- [path \(Protocols MPLS\) on page 280](#)
- [path-mtu on page 281](#)
- [policing \(Protocols MPLS\) on page 282](#)
- [policy-statement on page 283](#)
- [pop on page 287](#)
- [preference \(Protocols MPLS\) on page 288](#)
- [primary \(Protocols MPLS\) on page 289](#)

- [push](#) on page 290
- [record](#) on page 291
- [retry-limit](#) on page 292
- [revert-timer](#) on page 293
- [rsvp-error-hold-time](#) on page 294
- [secondary \(Protocols MPLS\)](#) on page 295
- [select](#) on page 296
- [signal-bandwidth](#) on page 296
- [smart-optimize-timer](#) on page 297
- [standby](#) on page 298
- [static-label-switched-path](#) on page 299
- [statistics \(Protocols MPLS\)](#) on page 301
- [swap](#) on page 302
- [switching-type](#) on page 303
- [te-class-matrix](#) on page 304
- [traffic-engineering \(Protocols MPLS\)](#) on page 305
- [transit \(Chained Composite Next Hops\)](#) on page 306
- [transit-lsp-association](#) on page 308
- [te-class-matrix](#) on page 309
- [to](#) on page 310
- [traceoptions \(Protocols MPLS\)](#) on page 311
- [traffic-engineering \(Protocols MPLS\)](#) on page 313
- [transit-lsp-association](#) on page 314

## adaptive

<b>Syntax</b>	<code>adaptive;</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 label-switched-path <i>lsp-name</i> ( <i>primary</i>   <i>secondary</i> ) <i>path-name</i> ], [edit protocols mpls label-switched-path <i>lsp-name</i> ], [edit protocols mpls label-switched-path <i>lsp-name</i> ( <i>primary</i>   <i>secondary</i> ) <i>path-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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</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. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Specify the bandwidth reallocation interval.
<b>Options</b>	<p><b><i>seconds</i></b>—Bandwidth reallocation interval, in seconds.</p> <p><b>Range:</b> 300 through 315,360,000 seconds</p> <p><b>Default:</b> 86,400 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 Automatic Bandwidth Allocation Interval on page 174</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. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<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 175</a></li> </ul>

## adjust-threshold-overflow-limit

<b>Syntax</b>	<code>adjust-threshold-overflow-limit <i>number</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 in Junos OS Release 7.5. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Specify the number of consecutive bandwidth overflow samples before triggering a bandwidth adjustment.
<b>Options</b>	<b><i>number</i></b> —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 175</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. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<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 175</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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>

## advertisement-hold-time

<b>Syntax</b>	<code>advertisement-hold-time <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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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><i>seconds</i></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>• <i>Damping Advertisement of LSP State Changes</i></li> </ul>

## associate-backup-pe-groups

<b>Syntax</b>	<code>associate-backup-pe-groups;</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 in Junos OS Release 9.0.
<b>Description</b>	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.
<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>• <i>Enabling Point-to-Point LSPs to Monitor Egress PE Routers</i></li> </ul>

## auto-bandwidth (MPLS Tunnel)

---


<b>Syntax</b>	<pre>auto-bandwidth {     adjust-interval <i>seconds</i>;     adjust-threshold <i>percent</i>;     adjust-threshold-activate-bandwidth <i>bps</i>     adjust-threshold-overflow-limit <i>number</i>;     adjust-threshold-underflow-limit <i>number</i>;     maximum-bandwidth <i>bps</i>;     minimum-bandwidth <i>bps</i>;     minimum-bandwidth-adjust-interval     minimum-bandwidth-adjust-threshold-change     minimum-bandwidth-adjust-threshold-value     monitor-bandwidth; }</pre>
<b>Hierarchy Level</b>	[edit protocols mpls label-switched-path <i>lsp-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<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 172</a></li><li>• <a href="#">request mpls lsp adjust-autobandwidth on page 332</a></li></ul>



## backup-pe-group

<b>Syntax</b>	<pre>backup-pe-group <i>group-name</i> {     backups [ <i>addresses</i> ];     local-address <i>address</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-options multicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> routing-options multicast],</p> <p>[edit routing-options multicast]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 12.3 for ACX Series routers.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.
<b>Options</b>	<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>
<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>• <i>Example: Configuring Ingress PE Redundancy</i></li> <li>• <i>Configuring Ingress PE Router Redundancy for Point-to-Multipoint LSPs</i></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>	<p>Statement introduced before Junos OS Release 7.4.          Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<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  <b>Default:</b> 0 (no bandwidth is reserved)</p>
<div>  <b>NOTE:</b> On the ACX Series, <i>bps</i> is the only supported option.         </div>	
	<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  <b>Default:</b> 0 (no bandwidth is reserved)</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.          routing-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Configuring Fast Reroute*
  - *Configuring the Bandwidth Value for LSPs*
  - *Configuring LSPs for DiffServ-Aware Traffic Engineering*
  - *Configuring Multiclass LSPs*

## 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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>	<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>• <i>Configuring Routers for DiffServ-Aware Traffic Engineering</i></li> </ul>

## bypass (Static LSP)

---

<b>Syntax</b>	<pre>bypass bypass-name {     bandwidth bps;     description string;     next-hop (address   interface-name   address/interface-name);     push out-label;     to address; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems logical-system-name protocols mpls static-label-switched-path     lsp-name], [edit protocols mpls static-label-switched-path lsp-name]</pre>
<b>Release Information</b>	Statement introduced before Junos OS Release 10.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	<p>Configure specific bandwidth and path constraints for a bypass ingress LSP. It is possible to configure multiple bypass LSPs individually. If you do not, they all share the same path and bandwidth constraints.</p> <p>The remaining statements are explained separately.</p>
<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>• <i>Configuring Static LSPs</i></li></ul>

## chained-composite-next-hop

**Syntax** chained-composite-next-hop {  
     ingress;  
     transit (Chained Composite Next Hops);  
 }

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-options forwarding-table],  
 [edit routing-options forwarding-table]



**NOTE:** The [edit logical-systems] hierarchy level is not supported on the QFX10000 switches.

**Release Information** Statement introduced in Junos OS Release 12.1.  
 Statement introduced in Junos OS Release 12.3 for ACX Series routers.  
 Statement introduced in Junos OS Release 15.1 for QFX10000 Series switches.

**Description** Allows you to configure the chained composite next hops for devices handling ingress or transit traffic in the network.

Chained composite next hops help to facilitate the handling of large volumes of transit traffic in the core of large networks by allowing the router to process much larger volumes of routes. A chained composite next hop allows the router to direct sets of routes sharing the same destination to a common forwarding next hop, rather than having each route also include the destination. In the event that a network destination is changed, rather than having to update all of the routes sharing that destination with the new information, just the shared forwarding next hop is updated with the new information. The chained composite next hops continue to point to this forwarding next hop which now contains the new destination.

On platforms containing only MPCs, such as PTX Series Packet Transport Routers, the MX80 router, the MX2020 router, and the QFX10000 switches, chained composite next hops are enabled by default. On MX Series 3D Universal Edge Routers containing both DPC and MPC FPCs and on T4000 Core Routers containing MPC and FPCs, chained composite next hops are disabled by default and need to be explicitly configured.



**NOTE:**

- Starting with Junos OS Release 13.3, for chained composite next hop feature to take effect for directly connected PE devices, the chassis must be configured to use the `enhanced-ip` option (in the case of MX Series 3D Universal Edge Routers containing both DPC and MPC FPCs) or the `enhanced-mode` option (in the case of T4000 Core Routers containing MPC and FPCs) in the network service mode, in addition to the `l3vpn` configuration.

For more information about configuring chassis network services, see the *Junos OS Administration Library for Routing Devices*.

- On MX Series routers, removing the chained-composite-next-hop statement from a PE device configuration causes all IBGP sessions to be torn down and triggers the BGP session to flap as well. A similar change on a router configured as a route reflector does not have any effect, however.

The following is a sample system log message that is generated to record such an event:

```
Nov  6 15:16:21.670 host PE1: rpd[6947]: bgp_peer_mgmt_clear:5995:
NOTIFICATION sent to 10.0.100.2 (External AS 100): code 6 (Cease)
subcode 4 (Administratively Reset), Reason: Management session cleared
BGP neighbor
```



**NOTE:** Starting in Junos OS Release 14.1, the `transit l3vpn` statement is enabled by default on PTX Series Packet Transport Routers only.

The remaining statements are explained separately.

**Default** This statement is disabled by default.

**Options** `ingress`—Enable or disable composite chained next hop for ingress traffic.  
`transit`—Enable or disable composite chained next hop for transit traffic.

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**

- *Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs*
- [Chained Composite Next Hops for Transit Devices for VPNs on page 166](#)
- *Example: Configuring Chained Composite Next Hops for Direct PE-PE Connections in VPNs*
- `ingress`  
[transit \(Chained Composite Next Hops\) on page 306](#)

## class-of-service (Protocols MPLS)

<b>Syntax</b>	<code>class-of-service cos-value;</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 static-label-switched-path <i>lsp-name</i> ingress],          [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>],          [edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.</p>
<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>• <i>Configuring Class of Service for MPLS LSPs</i></li> <li>• <i>Configuring the Ingress Router for Static LSPs</i></li> <li>• <i>Configuring Static LSPs</i></li> </ul>

## corouted-bidirectional

---

<b>Syntax</b>	corouted-bidirectional;
<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 12.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify that the label-switched path be established as a corouted bidirectional packet LSP. You cannot configure this statement at the same time as the <b>corouted-bidirectional-passive</b> statement.
<b>Default</b>	This statement 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>• <i>Configuring Corouted Bidirectional LSPs</i></li><li>• <a href="#">corouted-bidirectional-passive on page 240</a></li></ul>

## corouted-bidirectional-passive

---

<b>Syntax</b>	corouted-bidirectional-passive;
<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 12.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify that the label-switched path be a passive LSP associated with a bidirectional LSP when it is signaled at the ingress router. This passive LSP enables the MPLS application to utilize the reverse LSP. You cannot configure this statement at the same time as the <b>corouted-bidirectional</b> statement.
<b>Default</b>	This statement 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>• <i>Configuring Corouted Bidirectional LSPs</i></li><li>• <a href="#">corouted-bidirectional on page 240</a></li></ul>



## description (Protocols MPLS)

<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>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.</p>
<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>Configuring a Text Description for LSPs</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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>• <i>Configuring Routers for DiffServ-Aware Traffic Engineering</i></li></ul>

## disable (Protocols MPLS)

<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 <i>interface interface-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols mpls <i>label-switched-path lsp-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols mpls <i>label-switched-path lsp-name</i> auto-bandwidth], [edit protocols mpls], [edit protocols mpls <i>interface interface-name</i> ], [edit protocols mpls <i>label-switched-path lsp-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
<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><i>Minimum MPLS Configuration</i></li> </ul>

## exclude (for Fast Reroute)

---


<b>Syntax</b>	(exclude [ <i>group-names</i> ]   no-exclude);
<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. Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.
<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>	<b><i>group-names</i></b> —Names of one or more 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>• <i>Configuring Fast Reroute</i></li><li>• <i>admin-groups</i></li></ul>

## exclude-srlg

<b>Syntax</b>	exclude-srlg;
<b>Hierarchy Level</b>	<p>[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>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 11.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<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>	<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>• <i>Example: Excluding SRLG Links Completely for the Secondary LSP</i></li> </ul>

## explicit-null (Protocols MPLS)

---

<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. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
<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.
<div> <b>NOTE:</b> Junos OS does not support explicit null routes with next hops to virtual tunnel (vt-) interfaces.</div>	
<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>• <i>Configuring RSVP to Pop the Label on the Ultimate-Hop Router</i></li></ul>

## fast-reroute (Protocols MPLS)

<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>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.</p>
<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>	<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</a></li> <li>• <a href="#">Fast Reroute Overview on page 168</a></li> <li>• <a href="#">MPLS Feature Support on QFX Series and EX4600 Switches on page 134</a></li> <li>• <a href="#">Interprovider and Carrier-of-Carriers VPNs on page 165</a></li> </ul>

## forwarding-table

---

<b>Syntax</b>	<pre>forwarding-table {     export [ <i>policy--names</i> ];     (indirect-next-hop   no-indirect-next-hop); }</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 11.3 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure information about the routing device's forwarding table.  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>• <i>Configuring Per-Packet Load Balancing</i></li></ul>

## from (Protocols MPLS)

---

<b>Syntax</b>	<pre>from <i>address</i>;</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. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
<b>Description</b>	Specify the source address to use for the LSP.  The address you specify does not affect the outgoing interface used by the LSP.
<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>• <i>Configuring the Ingress and Egress Router Addresses for LSPs</i></li></ul>



## gpId

<b>Syntax</b>	<code>gpId (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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>• <i>Configuring MPLS LSPs for GMPLS</i></li> </ul>

## hop-limit

<b>Syntax</b>	<code>hop-limit <i>number</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 <a href="#">label-switched-path</a> <i>lsp-name</i>],  [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> <a href="#">fast-reroute</a>],  [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 logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],  [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],  [edit protocols mpls],  [edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i>],  [edit 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="#">primary</a>   <a href="#">secondary</a>) <i>path-name</i>],  [edit protocols rsvp interface <i>interface-name</i> link-protection],  [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	<p>Specify the maximum number of routers that an LSP can traverse. This limit can be applied to any of the following:</p> <ul style="list-style-type: none"> <li>LSPs—The configured hop limit includes the ingress and egress routers. You can specify a hop limit for an LSP and for both primary and secondary paths.</li> <li>Fast reroute detour—Specify the number of additional routers a fast reroute detour can traverse relative to the protected LSP. For example, if an LSP traverses 4 routers, any detour for the LSP can be no more than 10 router hops, including the ingress and egress routers.</li> <li>Link protection bypass—Specify the maximum number of routers that a link protection bypass can traverse.</li> </ul>
<b>Options</b>	<p><b><i>number</i></b>—Maximum number of hops.</p> <p><b>Range:</b> 2 through 255 (for an LSP or for a link protection bypass); 0 through 255 (for fast reroute)</p> <p><b>Default:</b> 255 (for an LSP or for a link protection bypass); 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</a></li> <li><a href="#">Limiting the Number of Hops in LSPs</a></li> <li><a href="#">Configuring Link Protection on Interfaces Used by LSPs</a></li> </ul>

## include-all (for Fast Reroute)

<b>Syntax</b>	(include-all [ <i>group-names</i> ]   no-include-all);
<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. Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.
<b>Description</b>	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>
<b>Options</b>	<b>group-names</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>• <i>Configuring Fast Reroute</i></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. Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.
<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>group-names</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>• <i>Configuring Fast Reroute</i></li> </ul>

## ingress (LSP)

<b>Syntax</b>	<pre> ingress {     bandwidth <i>bps</i>;     class-of-service <i>cos-value</i>;     description <i>string</i>;     entropy-label;     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>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i>],</p> <p>[edit protocols mpls static-label-switched-path <i>lsp-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 10.1.</p> <p><b>entropy-label</b> option introduced in Junos OS Release 14.1.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	<p>Configure an ingress LSR for a static LSP.</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>• <i>Configuring Static LSPs</i></li> </ul>

## install (Protocols MPLS)

<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. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
<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 inet.3 or inet6.3 routing table.
<b>Options</b>	<b>active</b> —(Optional) Install the route into the inet.0 or inet6.0 routing table. This allows you to issue a <b>ping</b> or <b>traceroute</b> command on this address.  <b>destination-prefix</b> —IPv4 or IPv6 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>• <i>Adding LSP-Related Routes to the inet.3 or inet6.3 Routing Table</i></li> </ul>

## interface (Protocols MPLS)

---

<b>Syntax</b>	<pre>interface (<i>interface-name</i>   all) {     disable;     admin-group [ <i>group-names</i> ];     srlg <i>srlg-name</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. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
<b>Description</b>	Enable MPLS on one or more interfaces.
<b>Options</b>	<p><b><i>interface-name</i></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 <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p> <p><b><i>srlg srlg-name</i></b>—Name of the SRLG to associate with an interface.</p> <p>The remaining options 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>• <i>Minimum MPLS Configuration</i></li><li>• <i>Configuring Static LSPs</i></li><li>• <i>Example: Configuring SRLG</i></li></ul>

## ipv6-tunneling

---

<b>Syntax</b>	ipv6-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. Statement introduced in Junos OS Release 14.1X53-D30 for QFX Series switches.
<b>Description</b>	Allow IPv6 routes to be resolved over an MPLS network by converting LDP and RSVP routes stored in the inet.3 routing table to IPv4-mapped IPv6 addresses and then copying them into the inet6.3 routing table. This routing table can be used to resolve next hops for both inet6 and inet6-vpn routes.
<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>• <i>Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks</i></li></ul>

## l2circuit

<b>Syntax</b>	<pre> l2circuit {   local-switching {     interface <i>interface-name</i> {       description <i>text</i>;     end-interface {       interface <i>interface-name</i>;       protect-interface <i>interface-name</i>;     }     ignore-mtu-mismatch;     protect-interface <i>interface-name</i>;   } } neighbor <i>address</i> {   interface <i>interface-name</i> {     backup-neighbor <i>address</i>;     bandwidth (<i>bandwidth</i>   <i>ctnumber bandwidth</i>);     community <i>community-name</i>;     connection-protection;     (control-word   no-control-word);     description <i>text</i>;     egress-protection;     encapsulation-type <i>type</i>;     ignore-encapsulation-mismatch;     ignore-mtu-mismatch;     mtu <i>mtu-number</i>;     protect-interface <i>interface-name</i>;     pseudowire-status-tlv hot-standby-vc-on;     psn-tunnel-endpoint <i>address</i>;     virtual-circuit-id <i>identifier</i>;   } } 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> &lt;flag-modifier&gt; &lt;disable&gt;; } </pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols], [edit protocols]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.</p>
<b>Description</b>	<p>Enables a Layer 2 circuit.</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>



- Related Documentation**
- *Configuring ATM Trunking on Layer 2 Circuits*
  - *Configuring Bandwidth Allocation and Call Admission Control in Layer 2 Circuits*
  - *Configuring Interfaces for Layer 2 Circuits*
  - *Configuring LDP for Layer 2 Circuits*
  - *Configuring Policies for Layer 2 Circuits*
  - *Configuring Static Layer 2 Circuits*
  - *Tracing Layer 2 Circuit Operations*

## label-switched-path (Protocols MPLS)

```

Syntax  label-switched-path lsp-name {
        disable;
        adaptive;
        admin-down;
        admin-group {
            exclude [ group-names ];
            include-all [ group-names ];
            include-any [ group-names ];
        }
        auto-bandwidth {
            adjust-interval seconds;
            adjust-threshold percentage;
            maximum-bandwidth bps;
            minimum-bandwidth bps;
            monitor-bandwidth;
        }
        bandwidth bps {
            ct0 bps;
            ct1 bps;
            ct2 bps;
            ct3 bps;
        }
        class-of-service cos-value;
        description text;
        entropy-label;
        fast-reroute {
            (bandwidth bps | bandwidth-percent percentage);
            (exclude [ group-names ] | no-exclude);
            hop-limit number;
            (include-all [ group-names ] | no-include-all);
            (include-any [ group-names ] | no-include-any);
        }
        from address;
        install {
            destination-prefix/prefix-length <active>;
        }
        inter-domain;
        ldp-tunneling;
        link-protection;
        lsp-attributes {
            encoding-type (ethernet | packet | pdh | sonet-sdh);
            gpipid (ethernet | hdlc | ipv4 | pos-scrambling-crc-16 | pos-no-scrambling-crc-16 |
                pos-scrambling-crc-32 | pos-no-scrambling-crc-32 | ppp);
            signal-bandwidth type;
            switching-type (fiber | lambda | psc-1 | tdm);
        }
        metric metric;
        no-cspf;
        no-decrement-ttl;
        node-link-protection;
        optimize-timer seconds;
        p2mp lsp-name;
    }

```

```

policing {
    filter filter-name;
    no-auto-policing;
}
preference preference;
primary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    priority setup-priority reservation-priority;
    (record | no-record);
    select (manual | unconditional);
    standby;
}
priority setup-priority reservation-priority;
(random | least-fill | most-fill);
(record | no-record);
retry-limit number;
retry-timer seconds;
revert-timer seconds;
secondary path-name {
    adaptive;
    admin-group {
        exclude [ group-names ];
        include-all [ group-names ];
        include-any [ group-names ];
    }
    bandwidth bps {
        ct0 bps;
        ct1 bps;
        ct2 bps;
        ct3 bps;
    }
    class-of-service cos-value;
    hop-limit number;
    no-cspf;
    no-decrement-ttl;
    optimize-timer seconds;
    preference preference;
    priority setup-priority reservation-priority;
    (record | no-record);
}

```

```

    select (manual | unconditional);
    standby;
}
soft-preemption;
standby;
to address;
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
}

```

<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>lsp-name</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>• <i>Minimum MPLS Configuration</i></li> <li>• <i>Configuring the Ingress and Egress Router Addresses for LSPs</i></li> <li>• <i>Configuring Primary and Secondary LSPs</i></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>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<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="#">Enabling LDP over RSVP-Established LSPs on page 34</a></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>• <i>Configuring Static LSPs</i></li> <li>• <i>Example: Configuring Point-to-Multipoint LSPs with Static Routes</i></li> </ul>

## log-updown (Protocols MPLS)

<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>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>The <b>mpls-lsp-traps</b> and <b>rfc-3812-traps</b> options added in Junos OS Release 9.0.</p> <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.</p>
<b>Description</b>	<p>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.</p>
<div style="display: flex; align-items: center;">  <p><b>NOTE:</b> System log messages for LSPs are generated by default. To disable the default logging of messages for LSPs, configure the <b>no-syslog</b> option under the <b>log-updown</b> statement.</p> </div>	
<b>Default</b>	<p>There is no default behavior for this statement. If you do not specify the options, the configuration cannot be committed.</p>
<b>Options</b>	<p><b>no-syslog</b>—Do not log a message to the system log file.</p> <p><b>no-trap</b>—Do not send an SNMP trap.</p> <p><b>syslog</b>—Log a message to the system log file.</p> <p><b>trap</b>—Send an SNMP trap.</p> <p><b>trap-path-down</b>—Send an SNMP trap when an LSP path goes down.</p> <p><b>trap-path-up</b>—Send an SNMP trap when an LSP path comes up.</p> <p>The <b>no-trap</b> statement is 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>• <i>Configuring System Log Messages and SNMP Traps for LSPs</i></li> <li>• <i>Network Management Administration Guide for Routing Devices</i></li> <li>• <a href="#">no-trap on page 272</a></li> <li>• <a href="#">traceoptions (Protocols MPLS) on page 311</a></li> </ul>

## **lsp-attributes**

<b>Syntax</b>	<pre>lsp-attributes {     encoding-type (ethernet   packet   pdh   sonet-sdh);     <a href="#">gp-id</a> (ethernet   hdlc   ipv4   pos-scrambling-crc-16   pos-no-scrambling-crc-16       pos-scrambling-crc-32   pos-no-scrambling-crc-32   ppp);     <a href="#">signal-bandwidth</a> type;     <a href="#">switching-type</a> (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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	<p>Define the parameters signaled during LSP setup. These usually determine the nature of the resource (label) allocated for the LSP.</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>• <i>Configuring MPLS LSPs for GMPLS</i></li> </ul>

## maximum-bandwidth (Protocols MPLS)

---

<b>Syntax</b>	<code>maximum-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. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<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 174</a></li></ul>

## metric (Protocols MPLS)

---

<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. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
<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 LSP Metrics</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. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<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 174</a></li> </ul>

## monitor-bandwidth

<b>Syntax</b>	<code>monitor-bandwidth;</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. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<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 177</a></li> </ul>

## mtu-signaling

---

<b>Syntax</b>	mtu-signaling;
<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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable MTU signaling in RSVP.
<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>• <i>Configuring MTU Signaling in RSVP</i></li></ul>

## no-cspf

<b>Syntax</b>	no-cspf;
<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>	<p>Statement introduced before Junos OS Release 7.4.          Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<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>	<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>• <i>Disabling Constrained-Path LSP Computation</i></li> <li>• <i>Configuring Explicit-Path LSPs</i></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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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</a></li><li>• <a href="#">no-propagate-ttl on page 270</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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Prevent the egress router address configured using the <b>to</b> statement from being installed into the inet.3 and inet.0 routing tables.
<b>Default</b>	The egress router address for an LSP is installed into the inet.3 and inet.0 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>• <i>Preventing the Addition of Egress Router Addresses to Routing Tables</i></li> <li>• <a href="#">to on page 310</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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	<p>Disable normal time-to-live (TTL) decrementing. You configure this statement once per router, and it affects all RSVP-signaled or LDP-signaled LSPs. When this router acts as an ingress router for an LSP, it pushes an MPLS header with a TTL value of 255, regardless of the IP packet TTL. When the router acts as the penultimate router, it pops the MPLS header without writing the MPLS TTL into the IP packet.</p> <p>When you add the <b>no-propagate-ttl</b> statement to the configuration or delete it from the configuration, the effect takes place immediately. There is no need to clear existing RSVP LSPs or LDP sessions.</p>
<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>• <i>Disabling Normal TTL Decrementing</i></li><li>• <i>Example: Disabling Normal TTL Decrementing in a VRF Routing Instance</i> (on <i>Layer 3 VPNs Feature Guide for Routing Devices</i> or in the <i>Junos VPNs Configuration Guide</i>)</li><li>• <a href="#">no-decrement-ttl on page 268</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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>Disabling Path Route Recording by LSPs</li> </ul>

## no-trap

---

<b>Syntax</b>	<pre>no-trap {     mpls-lsp-traps;     rfc-3812-traps; }</pre>
<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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>rfc-3812-traps</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>• <i>Configuring System Log Messages and SNMP Traps for LSPs</i></li><li>• <i>Network Management Administration Guide for Routing Devices</i></li><li>• <a href="#">traceoptions (Protocols MPLS) on page 311</a></li></ul>



## node-link-protection (Protocols MPLS)

<b>Syntax</b>	node-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. Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.
<b>Description</b>	Enable node and link protection on the specified LSP. To fully enable node and link protection, you also need to include the <b>link-protection</b> statement at the [edit protocols <b>rsvp</b> interface <i>interface-name</i> ] hierarchy level.
<b>Default</b>	Node and 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 Node Protection or Link Protection for LSPs</a></li> <li>• <a href="#">MPLS Feature Support on QFX Series and EX4600 Switches on page 134</a></li> <li>• <a href="#">Interprovider and Carrier-of-Carriers VPNs on page 165</a></li> </ul>

## oam (Protocols MPLS)

```

Syntax  oam {
        bfd-liveness-detection{
            failure-action teardown;
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            minimum-transmit-interval milliseconds;
            multiplier detection-time-multiplier;
        }
        lsp-ping-interval seconds;
        mpls-tp-mode;
        performance-monitoring {
            querier {
                loss {
                    traffic-class tc-value {
                        query-interval milliseconds;
                        measurement-quantity bytes|packets;
                        average-sample-size sample size;
                        loss-threshold loss threshold value;
                        loss-threshold-window number of samples for loss threshold;
                    }
                }
                delay {
                    traffic-class tc-value {
                        query-interval milliseconds;
                        padding-size size;
                        average-sample-size sample size;
                        rtt-delay-threshold rtt threshold value;
                        twcd-delay-threshold twcd threshold value;
                    }
                }
            }
            loss-delay {
                traffic-class tc-value {
                    query-interval milliseconds;
                    measurement-quantity bytes|packets;
                    padding-size size;
                    average-sample-size sample size;
                    loss-threshold loss threshold value;
                    loss-threshold-window number of samples for loss threshold;
                    rtt-delay-threshold rtt threshold value;
                    twcd-delay-threshold twcd threshold value;
                }
            }
        }
        responder {
            loss {
                min-query-interval milliseconds;
            }
            delay {
                min-query-interval milliseconds;
            }
        }
    }

```

}

<b>Hierarchy Level</b>	[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 <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. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric. <b>performance-monitoring</b> configuration statement introduced in Junos OS Release 15.1.
<b>Description</b>	Enable Operation, Administration, and Maintenance (OAM) for RSVP-signaled LSPs.
<b>Options</b>	<b>lsp-ping-interval seconds</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.  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>• <i>Configuring BFD for MPLS IPv4 LSPs</i></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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>• <i>Optimizing Signaled LSPs</i></li> </ul>

## optimize-hold-dead-delay

---

<b>Syntax</b>	<code>optimized-hold-dead-delay seconds;</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-switch-path <i>lsp-name</i> ], [edit protocols mpls], [edit protocols mpls label-switch-path <i>lsp-name</i> ]
<b>Description</b>	Allows you to specify the amount of time to delay the tear down of old paths after the router has switched traffic to new optimized paths. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The specified delay helps to ensure that old paths are not torn down before all routes have been switched over to the new optimized paths. This delay timer starts when the timer specified by the <b>optimize-switchover-dealy</b> statement has elapsed.
<b>Options</b>	<b>seconds</b> —Configure the time in seconds to wait before tearing down the old paths that were in use prior to the last LSP optimization. <b>Default:</b> 60 seconds <b>Range:</b> 0 through 65,535 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>• <i>Optimizing Signaled LSPs</i></li><li>• <a href="#">optimize-switchover-delay on page 277</a></li><li>• <a href="#">optimize-timer on page 278</a></li></ul>

## optimize-switchover-delay

<b>Syntax</b>	<code>optimize-switchover-delay <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 in Junos OS Release 11.1R1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Delays the switch over of LSPs to newly optimized paths. You only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers). The specified delay helps to ensure that the new optimized paths have been established before traffic is switched over from the old paths.
<b>Options</b>	<p><b><i>seconds</i></b>—Configure the time in seconds to wait before switching LSPs to newly optimized paths.</p> <p><b>Default:</b> 1 second</p> <p><b>Range:</b> 1 through 900 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>• <i>Optimizing Signaled LSPs</i></li> <li>• <a href="#">optimize-hold-dead-delay on page 276</a></li> <li>• <a href="#">optimize-timer on page 278</a></li> </ul>

## optimize-timer (Protocols MPLS)

<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],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols mpls <b>label-switched-path</b> <i>lsp-name</i>],</p> <p>[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>],</p> <p>[edit protocols mpls],</p> <p>[edit protocols mpls <b>label-switched-path</b> <i>lsp-name</i>],</p> <p>[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>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.</p>
<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 feature 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. Also, you only need to configure this statement on routers acting as the ingress for the affected LSPs (you do not need to configure this statement on transit or egress routers).</p> <p>To avoid extensive resource consumption that might result because of frequent path recomputations, or to avoid destabilizing the network as a result of constantly changing LSPs, we recommend that you either leave the timer value sufficiently large or disable the timer value.</p>
<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><i>Optimizing Signaled LSPs</i></li> </ul>

## p2mp (Protocols MPLS)

<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. Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.
<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>	<b><i>p2mp-lsp-name</i></b> —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>• <i>Configuring Primary and Branch LSPs for Point-to-Multipoint LSPs</i></li> </ul>

## path (Protocols MPLS)

<b>Syntax</b>	<pre>path <i>path-name</i> {     (<i>address</i>   <i>hostname</i>) &lt;strict   loose&gt;; }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series Virtual Chassis and Virtual Chassis Fabric.</p>
<b>Description</b>	<p>Create a named path and optionally specify the sequence of explicit routers that form the path.</p> <p>You must include this statement when configuring explicit LSPs.</p>
<b>Options</b>	<p><b>address</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>hostname</b>—See <b>address</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>loose</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>path-name</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>strict</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>• <i>Creating Named Paths</i></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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>• <i>Configuring MTU Signaling in RSVP</i></li></ul>

## policing (Protocols MPLS)

---

<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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>• <i>Configuring MPLS Firewall Filters and Policers</i></li><li>• <i>auto-policing</i></li></ul>

## policy-statement

<b>Syntax</b>	<pre> policy-statement <i>policy-name</i> {   term <i>term-name</i> {     from {       family <i>family-name</i>;       match-conditions;       policy <i>subroutine-policy-name</i>;       prefix-list <i>prefix-list-name</i>;       prefix-list-filter <i>prefix-list-name</i> match-type &lt;actions&gt;;       protocol <i>protocol-name</i>;       route-filter <i>destination-prefix</i> match-type &lt;actions&gt;;       source-address-filter <i>source-prefix</i> match-type &lt;actions&gt;;     }     to {       match-conditions;       policy <i>subroutine-policy-name</i>;     }     then <i>actions</i>;   }   then {     no-entropy-label-capability;   } } </pre>
<b>Hierarchy Level</b>	[edit dynamic policy-options], [edit logical-systems <i>logical-system-name</i> policy-options], [edit policy-options]
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5.</p> <p>Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>inet-mdt</b> option introduced in Junos OS Release 10.0R2.</p> <p>Statement introduced in Junos OS Release 11.3 for the QFX Series.</p> <p><b>route-target</b> option introduced in Junos OS Release 12.2.</p> <p>Statement introduced in Junos OS 14.1X53-D20 for the OCX Series.</p> <p><b>protocol</b> and <b>traffic-engineering</b> options introduced in Junos OS Release 14.2.</p> <p><b>no-entropy-label-capability</b> option introduced in Junos OS Release 15.1.</p>
<b>Description</b>	<p>Define a routing policy, including subroutine policies.</p> <p>A <i>term</i> is a named structure in which match conditions and actions are defined. Routing policies are made up of one or more terms. Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.</p> <p>Each term contains a set of match conditions and a set of actions:</p>

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

Generally, a router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of **accept** or **reject** is taken. If none of the terms in the policy match the route, the router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

If none of the match conditions of each term evaluates to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either **accept** or **reject**) that overrides any action intrinsic to the protocol.

The order of match conditions in a term is not relevant, because a route must match all match conditions in a term for an action to be taken.

To list the routing policies under the **[edit policy-options]** hierarchy level by **policy-statement *policy-name*** in alphabetical order, enter the **show policy-options** configuration command.

**Options** *actions*—(Optional) One or more actions to take if the conditions match. The actions are described in *Configuring Flow Control Actions*.

**family** *family-name*—(Optional) Specify an address family protocol. Specify **inet** for IPv4. Specify **inet6** for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify **iso**. For IPv4 multicast VPN traffic, specify **inet-mvpn**. For IPv6 multicast VPN traffic, specify **inet6-mvpn**. For multicast-distribution-tree (MDT) IPv4 traffic, specify **inet-mdt**. For BGP route target VPN traffic, specify **route-target**. For traffic engineering, specify **traffic-engineering**.



**NOTE:** When *family* is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the *family* statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

**from**—(Optional) Match a route based on its source address.

**match-conditions**—(Optional in **from** statement; required in **to** statement) One or more conditions to use to make a match. The qualifiers are described in *Routing Policy Match Conditions*.

**policy subroutine-policy-name**—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). Policy names cannot take the form **\_\_.\*-internal\_\_**, as this form is reserved. For information about how to configure subroutines, see *Understanding Policy Subroutines in Routing Policy Match Conditions*.

**no-entropy-label-capability**—(Optional) Disable the entropy label capability advertisement at egress or transit routes specified in the policy.

**policy subroutine-policy-name**—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" "). Policy names cannot take the form **\_\_.\*-internal\_\_**, as this form is reserved. For information about how to configure subroutines, see *Understanding Policy Subroutines in Routing Policy Match Conditions*.

**policy-name**—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").

**prefix-list prefix-list-name**—Name of a list of IPv4 or IPv6 prefixes.

**prefix-list-filter prefix-list-name**—Name of a prefix list to evaluate using qualifiers; *match-type* is the type of match (see *Configuring Prefix List Filters*), and *actions* is the action to take if the prefixes match.

**protocol** *protocol-name*—Name of the protocol used to control traffic engineering database import at the originating point.

**route-filter** *destination-prefix match-type <actions>*—(Optional) List of routes on which to perform an immediate match; *destination-prefix* is the IPv4 or IPv6 route prefix to match, *match-type* is the type of match (see *Configuring Route Lists*), and *actions* is the action to take if the *destination-prefix* matches.

**source-address-filter** *source-prefix match-type <actions>*—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. *source-prefix* is the IPv4 or IPv6 route prefix to match, *match-type* is the type of match (see *Configuring Route Lists*), and *actions* is the action to take if the *source-prefix* matches.

**term** *term-name*—Name that identifies the term. The term name must be unique in the policy. It can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" "). A policy statement can include multiple terms. We recommend that you name all terms. However, you do have the option to include an unnamed term which must be the final term in the policy. To configure an unnamed term, omit the **term** statement when defining match conditions and actions.

**to**—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

**then**—(Optional) Actions to take on matching routes. The actions are described in *Configuring Flow Control Actions* and *Configuring Actions That Manipulate Route Characteristics*.

<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>dynamic-db</i></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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>• <i>Configuring the Intermediate and Egress Routers for Static LSPs</i></li> <li>• <a href="#">swap on page 302</a></li> </ul>

## preference (Protocols MPLS)

<b>Syntax</b>	<code>preference <i>preference</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 logical-systems <i>logical-system-name</i> protocols mpls static-label-switched-path <i>lsp-name</i> ingress],  [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>],  [edit protocols mpls static-label-switched-path <i>lsp-name</i> ingress]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.  Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<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>preference</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.  routing-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Configuring Preference Values for LSPs</i></li> <li>• <i>Configuring Static LSPs</i></li> <li>• <i>Configuring Static LSPs</i></li> </ul>



## primary (Protocols MPLS)

<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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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 [edit mpls label-switched-path <i>lsp-name</i>] 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>Configuring Primary and Secondary LSPs</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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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 287</a></li><li>• <a href="#">swap on page 302</a></li><li>• <i>Configuring Static LSPs</i></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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>Disabling Path Route Recording by LSPs</li> </ul>

## retry-limit

---

<b>Syntax</b>	<code>retry-limit <i>number</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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>	<b><i>number</i></b> —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>• <i>Configuring the Connection Between Ingress and Egress Routers</i></li></ul>

## revert-timer

<b>Syntax</b>	<code>revert-timer seconds;</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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>seconds</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>	<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><i>Configuring Primary and Secondary LSPs</i></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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>	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>• <i>Improving Traffic Engineering Database Accuracy with RSVP PathErr Messages</i></li></ul>

## secondary (Protocols MPLS)

Syntax	<pre> secondary <i>path-name</i> {     adaptive;     admin-group {         exclude [ <i>group-names</i> ];         include-all [ <i>group-names</i> ];         include-any [ <i>group-names</i> ];     }     bandwidth <i>bps</i>;     class-of-service <i>cos-value</i>;     hop-limit <i>number</i>;     no-cspf;     no-decrement-ttl;     optimize-timer <i>seconds</i>;     preference <i>preference</i>;     priority <i>setup-priority reservation-priority</i>;     (record   no-record);     retry-limit <i>number</i>;     retry-timer <i>seconds</i>;     select (manual   unconditional);     standby; } </pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> ], [edit protocols mpls label-switched-path <i>lsp-name</i> ]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
Description	<p>Specify one or more secondary paths to use for the LSP. You can configure more than one secondary path. All secondary paths are equal, and the first one that is available is chosen.</p> <p>You can specify secondary paths even if you have not specified any primary paths.</p> <p>Optionally, you can specify preference, CoS, and bandwidth values for the secondary path, which override any equivalent values that you configure for the LSP (at the [edit mpls label-switched-path] hierarchy level).</p>
Options	<p><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>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <li>Configuring Primary and Secondary LSPs</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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>	<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).  <b>unconditional</b> —The path is always selected for carrying traffic, even if it is currently down or degraded (receiving errors).
<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>• <i>Configuring Primary and Secondary LSPs</i></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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>	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>• <i>Configuring MPLS LSPs for GMPLS</i></li> </ul>



## smart-optimize-timer

<b>Syntax</b>	<code>smart-optimize-timer 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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	<p>Enable the smart optimization timer. When you enable the smart optimization timer on a router, the Junos OS operates on the assumption that the original LSP path is preferable to any alternate or secondary path. When you enable the smart optimization timer and an LSP fails and its traffic is switched to an alternate path, the smart optimization timer starts and waits 3 minutes (this time is configurable). After 3 minutes have passed, the LSP is switched back to the original path. If the original path fails again and the LSP is switched to an alternate path again, the router waits 1 hour before attempting to switch the LSP back to its original path.</p> <p>If you want to disable the smart optimizer, you can set it to zero. The <b>smart-optimize-timer</b> value in seconds indicates the time before which the LSP is switched back to its primary path in case the primary path becomes available. Otherwise, the time to wait is controlled by the <b>optimize-timer</b>, which is usually set to a high value. Some ISPs have the <b>optimize-timer</b> set to once a day. Sometimes after the smart optimizer causes the LSP to be placed back on its primary path, the primary path goes down again within 60 minutes. When this happens, the <b>smart-optimize-timer</b> is disabled automatically, and the <b>optimize-timer</b> (regular path optimization) goes into effect. This is to protect against a flapping link being used.</p>
<b>Default</b>	The smart optimization timer is enabled by default.
<b>Options</b>	<p><b>seconds</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>• <i>Configuring the Smart Optimize Timer for LSPs</i></li> <li>• <i>Optimizing Signaled LSPs</i></li> <li>• <a href="#">optimize-aggressive on page 275</a></li> <li>• <a href="#">optimize-timer on page 278</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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>• <i>Configuring Hot Standby of Secondary Paths for LSPs</i></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.  
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Configure a static LSP.

**Options** *lsp-name*—Name of the path.

The remaining statements are explained separately.

<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Configuring Static LSPs</i></li></ul>

## statistics (Protocols MPLS)

<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>;   no-transit-statistics;   traffic-class-statistics;   transit-statistics-polling; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> protocols mpls], [edit protocols mpls]</pre>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p> <p><b>traffic-class-statistics</b> option introduced in Junos OS Release 14.2.</p>
<b>Description</b>	Enable MPLS statistics collection and reporting.
<b>Options</b>	<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 <code>mpls-stat</code> in the <code>/var/log</code> directory.</p> <p><b>files <i>number</i></b>—(Optional) Maximum number of trace files. When a trace file named <i>file</i> reaches its maximum size, it is renamed <i>file.0</i>, then <i>file.1</i>, and so on, until the maximum number of files is reached. Then, the oldest file is overwritten.</p> <p><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 <i>file</i> reaches this size, it is renamed <i>file.0</i>. When the <i>file</i> again reaches its maximum size, <i>file.0</i> is renamed <i>file.1</i> and <i>file</i> is renamed <i>file.0</i>. This renaming scheme continues until the maximum number of files is reached. Then the oldest trace file is overwritten.</p> <p>If you specify a maximum file size, you also must specify a maximum number of files with the <b>files</b> option.</p> <p><b>world-readable</b>—(Optional) Enable users to read the log file.</p> <p><b>Syntax:</b> Syntax: <b>xk</b> to specify KB, <b>xm</b> to specify MB, or <b>xg</b> to specify GB</p> <p><b>Range:</b> 10 KB through the maximum file size supported on your system</p>

**Default:** 1 MB

**traffic-class-statistics**—(Optional) Create counters that maintain data traffic statistics per traffic class at the ingress of all types of LSPs and egress of ultimate hop popping (UHP) point-to-point LSPs. These counters are not created by default and are required to be configured to perform traffic-class-scoped loss measurement.

**transit-statistics-polling**—(Optional) Enable the polling and display of MPLS statistics for LSPs transiting the router. By default, RSVP does not periodically poll for transit LSP statistics. You cannot configure this statement and the **no-transit-statistics** statement at the same time.

The remaining statements are explained separately.

**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring MPLS to Gather Statistics on page 187](#)
- [Configuring Automatic Bandwidth Allocation for LSPs on page 172](#)

## swap

<b>Syntax</b>	<code>swap 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> 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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	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.
<b>Options</b>	<b>out-label</b> —Manually assigned outgoing label value. <b>Range:</b> 0 through 1,048,575 <b>Default:</b> If you do not define the <b>out-label</b> option, the original label value remains unchanged.
<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 287</a></li> <li>• <a href="#">push on page 290</a></li> <li>• <a href="#">Configuring Static LSPs</a></li> </ul>

## switching-type

---

<b>Syntax</b>	<code>switching-type (fiber   lambda   psc-1   tdm);</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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the switching method for the LSP. The switching method can be one of the following values: <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>	<code>psc-1</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>• <i>Configuring MPLS LSPs for GMPLS</i></li> </ul>

## te-class-matrix

<b>Syntax</b>	<pre>te-class-matrix {     tenumber {         priority <i>priority</i>;         traffic-class {             <i>ctnumber</i> 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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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><i>ctnumber</i></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><i>priority priority</i></b>—Specify the priority of the class type. It can be one of eight values from 0 through 7.</p> <p><b><i>tenumber</i></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><i>traffic-class</i></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><i>Configuring Routers for DiffServ-Aware Traffic Engineering</i></li> </ul>



## traffic-engineering (Protocols MPLS)

<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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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 inet.3 routing table.</p> <p><b>bgp-igp</b>—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 routing table. If IGP shortcuts are enabled, the shortcut routes are automatically installed in the inet.0 routing table.</p> <p><b>bgp-igp-both-ribs</b>—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 and inet.3 routing tables. This option is used to support VPNs.</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>Configuring Traffic Engineering for LSPs</li> <li>Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)</li> </ul>

## transit (Chained Composite Next Hops)

**Syntax** transit {  
     (all | no-all);  
     (l2vpn | no-l2vpn);  
     (l3vpn | no-l3vpn);  
     (labeled-bgp | no-labeled-bgp);  
     (ldp | no-ldp);  
     (ldp-p2mp | no-ldp-p2mp);  
     lsp-statistics-from-route;  
     (rsvp | no-rsvp);  
     (rsvp-p2mp | no-rsvp-p2mp);  
     (static | no-static);  
 }

**Hierarchy Level** [edit logical-systems *logical-system-name* routing-options forwarding-table  
     chained-composite-next-hop],  
 [edit routing-options forwarding-table chained-composite-next-hop]



**NOTE:** The [edit logical-systems] hierarchy level is not supported on the QFX10000 switches.

**Release Information** Statement introduced in Junos OS Release 12.1.  
 Statement introduced in Junos OS Release 15.1 for QFX10000 Series switches.

**Description** Allows you to configure the chained composite next hops transit configuration options for devices handling transit traffic in the network. This statement and the associated functionality is available on PTX Packet Transport Routers and QFX10000 switches.

**Default** All of the **transit** statement options are enabled on PTX transport routers and QFX10000 switches. However, you can disable any of the statements with a **no** option.

**Options** **all | no-all**—Enable or disable chained composite next-hops for all of the possible packet transit protocols and applications. The **all | no-all** statements do not apply to the **lsp-statistics-from-route** statement.

**l2vpn | no-l2vpn**—Enable or disable chained composite next-hops for Layer 2 VPNs.

**l3vpn | no-l3vpn**—Enable or disable chained composite next-hops for Layer 3 VPNs.

**labeled-bgp | no-labeled-bgp**—Enable or disable chained composite next-hops for labeled BGP.

**ldp | no-ldp**—Enable or disable chained composite next-hops for LDP.

**ldp-p2mp | no-ldp-p2mp**—Enable or disable chained composite next-hops for LDP-signaled P2MP LSPs.

**lsp-statistics-from-route**—Enable LSP statistics collection from the route.

**rsvp | no-rsvp**—Enable or disable chained composite next-hops for RSVP.

**rsvp-p2mp | no-rsvp-p2mp**—Enable or disable chained composite next-hops for RSVP-signaled P2MP LSPs.

**static | no-static**—Chained composite next hops are enabled for transit static LSPs by default. You can also disable this functionality for transit static LSPs.



**NOTE:** On an MX series router with redundant Routing Engines and enhanced-ip mode configuration, enabling the **rsvp-p2mp** and **ldp-p2mp** statements under the **[edit routing-options forwarding-table chained-composite-next-hop transit]** hierarchy level causes ping from the current master logical system to fail at the time of a Routing Engine switchover.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**Related Documentation**

- *Accepting Route Updates with Unique Inner VPN Labels in Layer 3 VPNs*
- [chained-composite-next-hop on page 237](#)

## 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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>	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>• <i>Example: Configuring the MPLS Transport Profile for OAM</i></li></ul>

## te-class-matrix

<b>Syntax</b>	<pre>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 diffserv-te], [edit protocols mpls diffserv-te]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>tnumber</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>Configuring Routers for DiffServ-Aware Traffic Engineering</li> </ul>

to

---

<b>Syntax</b>	to <i>address</i> ;
<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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>• <i>Configuring the Ingress and Egress Router Addresses for LSPs</i></li></ul>

## traceoptions (Protocols MPLS)

<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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series. <b>ted-export</b> option introduced in Junos OS Release 14.2. <b>ted-import</b> option introduced in Junos OS Release 14.2. <b>lsp-history</b> option added in Junos OS Release 15.1.
<b>Description</b>	Configure MPLS tracing options at the protocol level or for a label-switched path.  To specify more than one tracing operation, include multiple <b>flag</b> statements.
<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>filename</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>files number</b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</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>flag</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>autobw-state</b>—Automatic bandwidth events.</li> <li>• <b>connection</b>—All circuit cross-connect (CCC) activity</li> <li>• <b>connection-detail</b>—Detailed CCC activity</li> </ul>

- **cspf**—CSPF computations
- **cspf-link**—Links visited during CSPF computations
- **cspf-node**—Nodes visited during CSPF computations
- **error**—MPLS error packets
- **graceful-restart**—Trace MPLS graceful restart events
- **lsp-history**—Trace LSP history events
- **lsping**—Trace lsping packets and return codes
- **nsr-synchronization**—Trace NSR synchronization events
- **nsr-synchronization-detail**—Trace NSR synchronization events in detail
- **state**—All LSP state transitions
- **static**—Trace static label-switched path
- **ted-export**—Trace leaking of entries from **lsdist.0** table into the traffic engineering database
- **ted-import**—Trace leaking traffic engineering database entries into the **lsdist.0** table
- **timer**—Timer usage

**no-world-readable**—(Optional) Allow only certain users to read the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

**world-readable**—(Optional) Allow any user to read the log file.

<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>• <i>Tracing MPLS and LSP Packets and Operations</i></li></ul>



## traffic-engineering (Protocols MPLS)

<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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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 inet.3 routing table.</p> <p><b>bgp-igp</b>—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 routing table. If IGP shortcuts are enabled, the shortcut routes are automatically installed in the inet.0 routing table.</p> <p><b>bgp-igp-both-ribs</b>—On both BGP and IGP destinations. Ingress routes are installed in the inet.0 and inet.3 routing tables. This option is used to support VPNs.</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>Configuring Traffic Engineering for LSPs</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. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<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>	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>• <i>Example: Configuring the MPLS Transport Profile for OAM</i></li></ul>


## CHAPTER 6

# Monitoring Commands for MPLS

- `clear mpls lsp`
- `monitor label-switched-path`
- `ping mpls bgp`
- `ping mpls l2circuit`
- `ping mpls l3vpn`
- `ping mpls lsp-end-point`
- `request mpls lsp adjust-autobandwidth`
- `show security keychain`
- `show link-management`
- `show link-management peer`
- `show link-management routing`
- `show link-management statistics`
- `show link-management te-link`
- `show mpls call-admission-control`
- `show mpls cspf`
- `show mpls diffserv-te`
- `show route forwarding-table`
- `show mpls interface`
- `show link-management statistics`
- `show link-management te-link`
- `show mpls call-admission-control`
- `show mpls cspf`
- `show mpls diffserv-te`
- `show route forwarding-table`
- `show mpls interface`
- `show mpls lsp`
- `show mpls lsp autobandwidth`
- `show mpls path`

- [show route table](#)
- [show route forwarding-table](#)
- [show mpls static-lsp](#)
- [show ted database](#)
- [show ted link](#)
- [show ted protocol](#)

## clear mpls lsp

<b>List of Syntax</b>	<a href="#">Syntax on page 317</a> <a href="#">Syntax (EX and QFX Series Switches) on page 317</a>
<b>Syntax</b>	<pre>clear mpls lsp &lt;autobandwidth&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;name <i>name</i>&gt; &lt;optimize   optimize-aggressive&gt; &lt;path <i>regular-expression</i>&gt; &lt;statistics&gt;</pre>
<b>Syntax (EX and QFX Series Switches)</b>	<pre>clear mpls lsp &lt;autobandwidth&gt; &lt;name <i>name</i>&gt; &lt;optimize   optimize-aggressive&gt; &lt;path <i>regular-expression</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.</p>
<b>Description</b>	Release the routes and states associated with MPLS label-switched paths (LSPs), and start new LSPs.
<div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p><b>CAUTION:</b> This command disconnects existing Resource Reservation Protocol (RSVP) sessions on the ingress routing device. If there is a time lag between the old path being torn down and the new path being set up, this command might impact traffic traveling along the LSPs.</p> </div> </div>	
<b>Options</b>	<p><b>none</b>—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.</p> <p><b>autobandwidth</b>—(Optional) Clear LSP autobandwidth counters.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>name <i>name</i></b>—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the <i>Junos Network Interfaces Configuration Guide</i>.</p> <p><b>optimize   optimize-aggressive</b>—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.</p>

**path *regular-expression***—(Optional) Clear the specific LSP path matching the specified regular expression.

**statistics**—(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (**name** and **path** options) on transit routers.

**Required Privilege Level**

clear

**Related Documentation**

- [show mpls lsp on page 386](#)
- [show rsvp session on page 575](#)

**List of Sample Output** [clear mpls lsp on page 318](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[clear mpls lsp](#)

```
user@host> clear mpls lsp
```

## monitor label-switched-path

**Syntax** `monitor label-switched-path lsp-name`  
`<logical-system (logical-system-name)>`

**Release Information** Command introduced before Junos OS Release 7.4.  
 Logical system support introduced in Junos OS Release 9.4.  
 Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.

**Description** Display the real-time status of the specified RSVP label-switched path (LSP). You can also use this command to monitor LSPs configured within logical systems.

**Options** `logical-system ( logical-system-name )`—(Optional) Perform this operation on all logical systems or on a particular logical system.

*lsp-name*—Name of the LSP.

**Additional Information** You can track the amount of traffic traversing an RSVP LSP and observe its essential parameters, such as uptime, ingress and egress addresses, labels, routes, and ports. Values are typically sampled every second. The display also allows you to scroll to other currently running LSPs. You cannot use this command to display information about static LSPs or LDP-signaled LSPs.

The output of this command shows how much each field has changed since you started the command or since you cleared the counters by using the `c` key. To control the output of the **monitor label-switched-path** command while it is running, use the keys listed in [Table 26 on page 319](#). The keys are not case-sensitive.

**Table 26: Output Control Keys for the monitor label-switched-path Command**

Key	Action
c	Clears the screen and refreshes the display for this LSP.
f	Freezes the display, preventing new information from being displayed.
l	Monitors a different LSP. After you type l, you can type the new LSP name.
n	Displays information about the next LSP (whose name is alphabetically higher than the current LSP name) configured on the router.
p	Goes to the previous LSP (whose name is alphabetically lower than the current LSP name) configured on the router.
q or Esc	Quits the command and returns to the command prompt.
t	Thaws, or restarts, the data display for this LSP.

**Required Privilege Level** trace

**List of Sample Output** [monitor label-switched-path on page 321](#)

**Output Fields** [Table 27 on page 320](#) describes the output fields for the **monitor label-switched-path** command. Output fields are listed in the approximate order in which they appear.

**Table 27: monitor label-switched-path Output Fields**

Field Name	Field Description
(1)	Displays the following information: <ul style="list-style-type: none"> <li>• <b>hostname</b>—Name of the router.</li> <li>• <b>Seconds</b>—Time elapsed since this display was started.</li> <li>• <b>Time</b>—Current local time.</li> </ul>
(2)	<b>Delay</b> —Length of the time delay, in milliseconds, required to obtain the information in the monitor display. The first number shows the current sampling delay. The second number shows the shortest delay recorded to date. The third number shows the worst delay recorded to date. This delay can vary substantially depending on the system load.
(3)	Displays the following: <ul style="list-style-type: none"> <li>• <b>To</b>—Destination address of the LSP.</li> <li>• <b>From</b>—Originating address of the LSP.</li> <li>• <b>State</b>—Current state of the LSP: <b>Up</b> or <b>Down</b>.</li> </ul>
(4)	Displays the following: <ul style="list-style-type: none"> <li>• <b>LSPName</b>—Name of the LSP.</li> <li>• <b>Type</b>—Type of LSP: <b>Ingress</b>, <b>Egress</b>, or <b>Transit</b>.</li> </ul>
(5)	Displays the following: <ul style="list-style-type: none"> <li>• <b>Label in</b>—Incoming label of the LSP.</li> <li>• <b>Label out</b>—Outgoing label of the LSP.</li> </ul>
(6)	<b>Port number</b> —Port number for the sending router, the port number for the receiving router, and the protocol ID. For MPLS traffic engineering applications, the protocol ID is always 0.
(7/8)	<b>Record route</b> —All intermediate and egress router addresses for this LSP.
(9/10/11)	Displays traffic statistics: <ul style="list-style-type: none"> <li>• <b>Output packets</b>—Number of packets that have traversed this LSP, and the change (delta) in the number since the last sample, typically 1 second ago.</li> <li>• <b>Output bytes</b>—Number of bytes that have traversed this LSP, and the change (delta) in the number since the last sample, typically 1 second ago.</li> </ul>
(12)	Displays any errors the router encountered while attempting to retrieve information on the LSP.
(13)	Lists the keyboard commands you can use to navigate to other LSPs. For a description of the keyboard commands, see <a href="#">Table 26 on page 319</a> .



## Sample Output

### monitor label-switched-path

```
user@host> monitor label-switched-path
(1) host                               Seconds: 112           Time: 15:32:22
(2)                                     Delay: 0/0/0
(3) To 10.10.10.16, From 10.10.10.17, state: Up
(4)  LSPname: k, type: Ingress
(5)  Label in: -, Label out: 126000
(6)  Port number: sender 1, receiver 45583, protocol 0
(7)  Record Route: <self> 192.168.224.196
(8)    192.168.224.202 192.168.224.179
(9)  Traffic statistics:
(10)    Output packets:                0                      Current delta [0]
(11)    Output bytes:                  0                      [0]
(12)
(13)Next='n', Prev='p', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c',
    LSP='l'
```

## ping mpls bgp

**Syntax** ping mpls bgp fec  
 <bottom-label-ttl>  
 <count *count*>  
 <destination *address*>  
 <detail>  
 <exp *forwarding-class*>  
 <instance *routing-instance-name*>  
 <logical-system (all | *logical-system-name*)>  
 <size *bytes*>  
 <source *source-address*>  
 <sweep>

**Release Information** Command introduced in Junos OS Release 11.1.

**Description** Check the operability of MPLS BGP-signaled label-switched path (LSP) connections. Press Ctrl+c to interrupt a **ping mpls bgp** command.



**NOTE:** The **ping mpls bgp fec** command only supports single paths.

- Options**
- bottom-label-ttl**—(Optional) Time-to-live (TTL) value for the bottom label in the label stack. The range of values is 1 through 255. The default value is **255**.
  - count *count***—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is **5**.
  - destination *address***—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.
  - detail**—(Optional) Display detailed information about the echo requests sent and received.
  - exp *forwarding-class***—(Optional) Value of the forwarding class for the MPLS ping packets.
  - fec**—Ping a BGP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.
  - instance *routing-instance-name***—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.
  - logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on the specified logical system.
  - size *bytes***—(Optional) Size of the LSP ping request packet (88 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.

**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only BGP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls bgp fec count on page 323](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately. To display the error codes, use the **detail** option (for example, **ping mpls bgp 10.255.245.222 detail**).

## Sample Output

### ping mpls bgp fec count

```
user@host> ping mpls bgp 10.255.245.222 count 10
!!!xxx...x--- 1sping statistics ---10 packets transmitted, 3 packets received,
70% packet loss 4 packets received with error status, not counted as received.
```

## ping mpls l2circuit

<b>Syntax</b>	<p>ping mpls l2circuit (interface <i>interface-name</i>   virtual-circuit <i>virtual-circuit-id</i> neighbor <i>address</i>)</p> <p>&lt;count <i>count</i>&gt;</p> <p>&lt;destination <i>address</i>&gt;</p> <p>&lt;detail&gt;</p> <p>&lt;exp <i>forwarding-class</i>&gt;</p> <p>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</p> <p>reply-mode (application-level-control-channel   ip-udp   no-reply)</p> <p>&lt;size <i>bytes</i>&gt;</p> <p>&lt;source <i>source-address</i>&gt;</p> <p>&lt;sweep&gt;</p> <p>&lt;v1&gt;</p>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p> <p>The <b>size</b> and <b>sweep</b> options were introduced in Junos OS Release 9.6.</p> <p>The <b>reply-mode</b> option and its suboptions are introduced in Junos OS Release 10.4R1.</p>
<b>Description</b>	<p>Check the operability of the MPLS Layer 2 circuit connections. Type Ctrl+c to interrupt a ping mpls l2circuit command.</p>
<b>Options</b>	<p><b>count</b> <i>count</i>—(Optional) Number of ping requests to send. If <b>count</b> is not specified, five ping requests are sent. The range of values is 1 through <b>1,000,000</b>. The default value is 5.</p> <p><b>destination</b> <i>address</i>—(Optional) Specify an address other than the default (<b>127.0.0.1/32</b>) for the ping echo requests. The address can be anything within the <b>127/8</b> subnet.</p> <p><b>detail</b>—(Optional) Display detailed information about the echo requests sent and received.</p> <p><b>exp</b> <i>forwarding-class</i>—(Optional) Value of the forwarding class for the MPLS ping packets.</p> <p><b>interface</b> <i>interface-name</i>—Ping an interface configured for the Layer 2 circuit on the egress provider edge (PE) router.</p> <p><b>logical-system</b> (all   <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on the specified logical system.</p> <p><b>reply-mode</b>—(Optional) Reply mode for the ping request. This option has the following suboptions:</p> <p><b>application-level-control-channel</b>—Reply using an application level control channel.</p> <p><b>ip-udp</b>—Reply using an IPv4 or IPv6 UDP packet.</p> <p><b>no-reply</b>—Do not reply to the ping request.</p>



**NOTE:** The **reply-mode** option and its suboptions **application-level-control-channel**, **ip-udp**, and **no-reply** are also available in Junos OS Release 10.2R4 and 10.3R2.

**size bytes**—(Optional) Size of the label-switched path (LSP) ping request packet (96 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

**source source-address**—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**vt**—(Optional) Use the type 9 Layer 2 circuit type, length, and value (TLV).

**virtual-circuit virtual-circuit-id neighbor address**—Ping the virtual circuit identifier on the egress PE router or switch and the specified neighbor, testing the integrity of the Layer 2 circuit between the ingress and egress PE routers or switches.

**Additional Information** You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls l2circuit interface on page 326](#)  
[ping mpls l2circuit virtual-circuit detail on page 326](#)  
[ping mpls l2circuit interface <interface-name> reply-mode on page 326](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping mpls l2circuit interface

```
user@host> ping mpls l2circuit interface so-1/0/0.1
Request for seq 1, to interface 69, labels <100000, 100208>, packet size 100
Reply for seq 1, return code: Egress-ok, time: 0.439 ms
```

### ping mpls l2circuit virtual-circuit detail

```
user@host> ping mpls l2circuit virtual-circuit 200 neighbor 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100048, 100128>, packet size 100

Reply for seq 1, return code: Egress-ok time: 0.539 ms
```

### ping mpls l2circuit interface <interface-name> reply-mode

```
user@host> ping mpls l2circuit interface lt-1/2/0.21 reply-mode application-level-control-channel
!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

## ping mpls l3vpn

<b>Syntax</b>	<pre>ping mpls l3vpn prefix <i>prefix-name</i> &lt;l3vpn-name&gt; &lt;bottom-label-ttl&gt; &lt;count <i>count</i>&gt; &lt;destination <i>address</i>&gt; &lt;detail&gt; &lt;exp <i>forwarding-class</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;size <i>bytes</i>&gt; &lt;source <i>source-address</i>&gt; &lt;sweep&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>The <b>size</b> and <b>sweep</b> options were introduced in Junos OS Release 9.6.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	<p>Check the operability of a MPLS Layer 3 virtual private network (VPN) connection. Type Ctrl+c to interrupt a <b>ping mpls l3vpn</b> command.</p>
<b>Options</b>	<p><b>bottom-label-ttl</b>—(Optional) Display the time-to-live value for the bottom label in the label stack.</p> <p><b>count <i>count</i></b>—(Optional) Number of ping requests to send. If <b>count</b> is not specified, five ping requests are sent. The range of values is 1 through <b>1,000,000</b>. The default value is <b>5</b>.</p> <p><b>destination <i>address</i></b>—(Optional) Specify an address other than the default (<b>127.0.0.1/32</b>) for the ping echo requests. The address can be anything within the <b>127/8</b> subnet.</p> <p><b>detail</b>—(Optional) Display detailed information about the echo requests sent and received.</p> <p><b>exp <i>forwarding-class</i></b>—(Optional) Value of the forwarding class for the MPLS ping packets.</p> <p><b><i>l3vpn-name</i></b>—(Optional) Layer 3 VPN name.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on the specified logical system.</p> <p><b>prefix <i>prefix-name</i></b>—Ping to test whether a prefix is present in a provider edge (PE) router's or switch's VPN routing and forwarding (VRF) table, by means of a Layer 3 VPN destination prefix. This option does not test the connection between a PE router or switch and a customer edge (CE) router or switch.</p> <p><b>size <i>bytes</i></b>—(Optional) Size of the label-switched path (LSP) ping request packet (<b>96</b> through <b>65468</b> bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.</p>

**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

If the Layer 3 VPN traffic transits a route reflector within the network, the **ping mpls l3vpn** command does not work.

**Required Privilege Level** network

**List of Sample Output** [ping mpls l3vpn on page 328](#)  
[ping mpls l3vpn detail on page 328](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code these packets are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping mpls l3vpn

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32
!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

### ping mpls l3vpn detail

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100128, 100112>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <100128, 100112>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <100128, 100112>
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <100128, 100112>
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 68, labels <100128, 100112>
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```





## ping mpls lsp-end-point

<b>Syntax</b>	<pre>ping mpls lsp-end-point <i>prefix-name</i> &lt;count <i>count</i>&gt; &lt;destination <i>address</i>&gt; &lt;detail&gt; &lt;exp <i>forwarding-class</i>&gt; &lt;instance <i>routing-instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;size <i>bytes</i>&gt; &lt;source <i>source-address</i>&gt; &lt;sweep&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>The <b>size</b> and <b>sweep</b> options were introduced in Junos OS Release 9.6.</p> <p>The <b>instance</b> option was introduced in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	<p>Check the operability of MPLS label-switched path (LSP) endpoint connections. Type Ctrl+c to interrupt a <b>ping mpls</b> command.</p>
<b>Options</b>	<p><b>count</b> <i>count</i>—(Optional) Number of ping requests to send. If <b>count</b> is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.</p> <p><b>destination</b> <i>address</i>—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.</p> <p><b>detail</b>—(Optional) Display detailed information about the echo requests sent and received.</p> <p><b>exp</b> <i>forwarding-class</i>—(Optional) Value of the forwarding class for the MPLS ping packets.</p> <p><b>instance</b> <i>routing-instance-name</i>—(Optional) Ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP connection.</p> <p><b>logical-system</b> (all   <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on the specified logical system.</p> <p><b>prefix-name</b>—LDP forwarding equivalence class (FEC) prefix or RSVP LSP endpoint address.</p> <p><b>size</b> <i>bytes</i>—(Optional) Size of the LSP ping request packet. If the endpoint is LDP-based, the minimum size of the packet is 88 bytes. If the endpoint is RSVP-based, the minimum size of the packet is 100 bytes. The maximum size in either case is 65468 bytes.</p> <p><b>source</b> <i>source-address</i>—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).</p>

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls lsp-end-point detail on page 331](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

## Sample Output

### [ping mpls lsp-end-point detail](#)

```
user@host> ping mpls lsp-end-point 10.255.245.119 detail
Route to end point address is via LDP FEC
Request for seq 1, to interface 67, label 100032
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, label 100032
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 67, label 100032
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 67, label 100032
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 67, label 100032
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

## request mpls lsp adjust-autobandwidth

---

<b>List of Syntax</b>	<a href="#">Syntax on page 332</a> <a href="#">Syntax (EX and QFX Series Switches) on page 332</a>
<b>Syntax</b>	<code>request mpls lsp adjust-autobandwidth</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code> <code>&lt;name <i>lsp-name</i>&gt;</code>
<b>Syntax (EX and QFX Series Switches)</b>	<code>request mpls lsp adjust-autobandwidth</code> <code>&lt;name <i>lsp-name</i>&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	<p>Manually trigger a bandwidth allocation adjustment for active label-switched paths (LSPs).</p> <p>Without running this command, the bandwidth adjustment is recomputed at a configurable interval. The default interval is 5 minutes. If you do not want to wait for the periodic adjustment (for example, during a software demonstration), this command is useful.</p> <p>During bandwidth allocation adjustment, the LSP stays up to enable the bandwidth to be changed without dropping any traffic. This functionality is often referred to as <i>make-before-break</i>.</p>
<b>Options</b>	<p><b>none</b>—Manually trigger a bandwidth allocation adjustment for all active LSP paths.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>name <i>lsp-name</i></b>—(Optional) Manually trigger a bandwidth allocation adjustment on the specified LSP only.</p>
<b>Additional Information</b>	<p>For this command to work properly, the following conditions must exist:</p> <ul style="list-style-type: none"><li>• Automatic bandwidth allocation must be enabled on the LSP. The parameters for adjustment interval and maximum average bandwidth are not reset after you issue the <b>request mpls lsp adjust-autobandwidth</b> command.</li><li>• The difference between the adjusted bandwidth and the current LSP path bandwidth must be greater than the threshold limit.</li></ul>
<b>Required Privilege Level</b>	clear, maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">auto-bandwidth on page 232</a></li><li>• <a href="#">Configuring Automatic Bandwidth Allocation for LSPs on page 172</a></li></ul>

**List of Sample Output** [request mpls lsp adjust-auto-bandwidth on page 333](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[request mpls lsp adjust-auto-bandwidth](#)

```
user@host> request mpls lsp adjust-auto-bandwidth
```

## show security keychain

<b>Syntax</b>	show security keychain <brief   detail>
<b>Release Information</b>	Command introduced in Junos OS Release 11.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Display information about authentication keychains configured for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, the Bidirectional Forwarding Detection (BFD) protocol, and the Intermediate System-to-Intermediate System (IS-IS) protocol.
<b>Options</b>	<b>none</b> —Display information about authentication keychains.  <b>brief   detail</b> —(Optional) Display the specified level of output.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show security keychain brief on page 335</a> <a href="#">show security keychain detail on page 336</a>
<b>Output Fields</b>	<a href="#">Table 28 on page 334</a> describes the output fields for the <b>show security keychain</b> command. Output fields are listed in the approximate order in which they appear.

**Table 28: show security keychain Output Fields**

Field Name	Field Description	Level of Output
<b>keychain</b>	The name of the keychain in operation.	All levels
<b>Active-ID Send</b>	Number of routing protocols packets sent with the active key.	All levels
<b>Active-ID Receive</b>	Number of routing protocols packets received with the active key.	All levels
<b>Next-ID Send</b>	Number of routing protocols packets sent with the next key.	All levels
<b>Next-ID Receive</b>	Number of routing protocols packets received with the next key.	All levels
<b>Transition</b>	Amount of time until the current key will be replaced with the next key in the keychain.	All levels
<b>Tolerance</b>	Configured clock-skew tolerance, in seconds, for accepting keys for a key chain.	All levels
<b>Id</b>	Identification number configured for the current key.	<b>detail</b>
<b>Algorithm</b>	Authentication algorithm configured for the current key.	<b>detail</b>

Table 28: show security keychain Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>State</b>	<p>State of the current key.</p> <p>The value can be:</p> <ul style="list-style-type: none"> <li>• <b>receive</b></li> <li>• <b>send</b></li> <li>• <b>send-receive</b></li> </ul> <p>For the active key, the <b>State</b> can be <b>send-receive</b>, <b>send</b>, or <b>receive</b>. For keys that have a future start time, the <b>State</b> is <b>inactive</b>. Compare the <b>State</b> field to the <b>Mode</b> field.</p>	<b>detail</b>
<b>Option</b>	<p>For IS-IS only, the option determines how Junos OS encodes the message authentication code in routing protocol packets.</p> <p>The values can be:</p> <ul style="list-style-type: none"> <li>• <b>basic</b>—Based on RFC 5304.</li> <li>• <b>isis-enhanced</b>—Based on RFC 5310.</li> </ul> <p>The default value is <b>basic</b>. When you configure the <b>isis-enhanced</b> option, Junos OS sends RFC 5310-encoded routing protocol packets and accepts both RFC 5304-encoded and RFC 5310-encoded routing protocol packets that are received from other devices.</p> <p>When you configure <b>basic</b> (or do not include the <b>options</b> statement in the key configuration) Junos OS sends and receives RFC 5304-encoded routing protocols packets, and drops 5310-encoded routing protocol packets that are received from other devices.</p> <p>Because this setting is for IS-IS only, the TCP and the BFD protocol ignore the encoding option configured in the key.</p>	<b>detail</b>
<b>Start-time</b>	Time that the current key became active.	<b>detail</b>
<b>Mode</b>	<p>Mode of each key (Informational only.)</p> <p>The value can be</p> <ul style="list-style-type: none"> <li>• <b>receive</b></li> <li>• <b>send</b></li> <li>• <b>send-receive</b></li> </ul> <p>The mode of the key is based on the configuration. Suppose you configure two keys, one with a start-time of today and the other with a start-time of next week. For both keys, the <b>Mode</b> can be <b>send-receive</b>, <b>send</b>, or <b>receive</b>, regardless of the configured start-time. Compare the <b>Mode</b> field to the <b>State</b> field.</p>	<b>detail</b>

## Sample Output

show security keychain brief

```
user@host> show security keychain brief
```

keychain	Active-ID		Next-ID		Transition	Tolerance
	Send	Receive	Send	Receive		
hakr	3	3	1	1	1d 23:58	3600

#### show security keychain detail

```
user@host> show security keychain detail
keychain              Active-ID      Next-ID      Transition  Tolerance
                      Send  Receive    Send  Receive
hakr                  3      3          1      1          1d 23:58    3600
  Id 3, Algorithm hmac-md5, State send-receive, Option basic
  Start-time Wed Aug 11 16:28:00 2010, Mode send-receive
  Id 1, Algorithm hmac-md5, State inactive, Option basic
  Start-time Fri Aug 20 11:30:57 2010, Mode send-receive
```



## show link-management

<b>Syntax</b>	show link-management
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) peer and traffic engineering link information.
<b>Options</b>	This command has no options.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show link-management peer on page 341</a></li> <li>• <a href="#">show link-management routing on page 343</a></li> <li>• <a href="#">show link-management statistics on page 346</a></li> <li>• <a href="#">show link-management te-link on page 348</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show link-management on page 340</a>
<b>Output Fields</b>	Table 29 on page 337 describes the output fields for the <b>show link-management</b> command. Output fields are listed in the approximate order in which they appear.

**Table 29: show link-management Output Fields**

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is 0 through 64,000.
State	State of the peer: <b>Up</b> or <b>Down</b> .
Control address	Address to which a control channel is established.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296.
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296.
State	State of the control channel: <b>Up</b> or <b>Down</b> .
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is 1 through 4,294,967,295.
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is 0 through 4,294,967,295.

Table 29: show link-management Output Fields (*continued*)

Field Name	Field Description
<b>Flags</b>	Code that provides information about the control channel. Currently supports only code value <b>R</b> , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
<b>TE links</b>	Traffic-engineered links that are managed by their peer.
<b>TE link name</b>	Name of the traffic-engineered link.
<b>State</b>	State of the traffic-engineered link: <b>Up</b> , <b>Down</b> , or <b>Init</b> .
<b>Local identifier</b>	Identifier of the local side of the link.
<b>Remote identifier</b>	Identifier of the remote side of the link.
<b>Local address</b>	Address of the local side of the link.
<b>Remote address</b>	Address of the remote side of the link.
<b>Encoding</b>	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b> , <b>Ethernet</b> , <b>Packet</b> , and <b>PDH</b> .
<b>Switching</b>	Type of switching that can be performed on the traffic-engineered link. Supported values are <b>PSC-1</b> and <b>Packet</b> .
<b>Minimum bandwidth</b>	Smallest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link (in bps).
<b>Maximum bandwidth</b>	Largest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
<b>Total bandwidth</b>	Sum of the bandwidth, in bits per second (bps) and megabits per second (Mbps), of all interfaces that are members of the link.
<b>Available bandwidth</b>	Sum of the bandwidths of all interfaces that are members of the link and that are not yet allocated (in bps).
<b>Name</b>	Name of the interface.
<b>State</b>	State of the interface: <b>Up</b> or <b>Down</b> .
<b>Local ID</b>	Identifier of the local side of the interface.
<b>Remote ID</b>	Identifier of the remote side of the interface.
<b>Bandwidth</b>	Bandwidth, in bps or Mbps, of the member interface.
<b>Used</b>	Whether the resource is allocated to an LSP: <b>Yes</b> or <b>No</b> .

Table 29: show link-management Output Fields (*continued*)

Field Name	Field Description
LSP-name	LSP name.

## Sample Output

### show link-management

```
user@host> show link-management
Peer name: PEER-A, System identifier: 11973
State: Up, Control address: 10.255.245.4
  CC local ID CC remote ID State      TxSeqNum  RcvSeqNum  Flags
    24547      24547 Up          1027      1026
TE links:
  pro4-ba

TE link name: pro4-ba, State: Init
Local identifier: 2662, Remote identifier: 0, Encoding: SDH/SONET, Switching:
PSC-1,
Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:
155.52Mbps,
Available bandwidth: 155.52Mbps
  Name          State Local ID Remote ID    Bandwidth Used  LSP-name
  so-1/0/2      Up          21271      0          155.52Mbps    No
```

## show link-management peer

<b>Syntax</b>	show link-management peer <name <i>peer-name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) peer link information.
<b>Options</b>	<b>none</b> —Display all peer link information.  <b>name <i>peer-name</i></b> —(Optional) Display information for the specified peer only.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 337</a></li> <li>• <a href="#">show link-management routing on page 343</a></li> <li>• <a href="#">show link-management statistics on page 346</a></li> <li>• <a href="#">show link-management te-link on page 348</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show link-management peer on page 342</a>
<b>Output Fields</b>	Table 30 on page 341 describes the output fields for the <b>show link-management peer</b> command. Output fields are listed in the approximate order in which they appear.

**Table 30: show link-management peer Output Fields**

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is 0 through 64,000.
State	State of the peer: <b>Up</b> or <b>Down</b> .
Control address	Address to which a control channel is established.
Hello interval	How often the routing device sends Link Management Protocol (LMP) hello packets.
Hello dead interval	How long LMP waits before declaring the control channel to be dead. This is an interval during which the routing device receives no LMP hello packets from the neighbor on a control that is active or up.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is 1 through 4,294,967,296.
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is 1 through 4,294,967,296.

Table 30: show link-management peer Output Fields (*continued*)

Field Name	Field Description
<b>State</b>	State of the control channel: <b>Up</b> or <b>Down</b> .
<b>TxSeqNum</b>	Sequence number of the hello message being sent to the peer. The range of values is <b>1</b> through <b>4,294,967,295</b> .
<b>RcvSeqNum</b>	Sequence number of the last hello message received from the peer. The range of values is <b>0</b> through <b>4,294,967,295</b> .
<b>Flags</b>	Code that provides information about the control channel. Currently supports only code value <b>R</b> , which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
<b>TE links</b>	Traffic-engineered links that are managed by their peer.

## Sample Output

### show link-management peer

```

user@host> show link-management peer
Peer name: sonet, System identifier: 41448
State: Up, Control address: 70.70.70.70
Hello interval: 10000, Hello dead interval: 30000
  CC local ID CC remote ID State      TxSeqNum  RcvSeqNum  Flags
    3265           0 ConfSnd         1          0 R
TE links:
  to-sonet

```

## show link-management routing

<b>Syntax</b>	show link-management routing <peer <name <i>name</i> >   te-link <name <i>name</i> >> <resource <name <i>name</i> >>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) peer or traffic engineering link information from the routing process.
<b>Options</b>	<p><b>none</b>—Display all peer and traffic-engineered link information.</p> <p><b>peer &lt;name <i>name</i>&gt;</b>—(Optional) Display information for all peers or for the specified peer only.</p> <p><b>resource &lt;name <i>name</i>&gt;</b>—(Optional) Display information for all resources or for the specified resource only.</p> <p><b>te-link &lt;name <i>name</i>&gt;</b>—(Optional) Display information for all traffic-engineered forwarding paths or for the specified path only.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 337</a></li> <li>• <a href="#">show link-management peer on page 341</a></li> <li>• <a href="#">show link-management statistics on page 346</a></li> <li>• <a href="#">show link-management te-link on page 348</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show link-management routing on page 345</a>
<b>Output Fields</b>	Table 31 on page 343 describes the output fields for the <b>show link-management routing</b> command. Output fields are listed in the approximate order in which they appear.

**Table 31: show link-management routing Output Fields**

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is 0 through 64,000.
State	State of the peer: Up or Down.
Control address	Address to which a control channel is established.
Control channel	Interface over which control packets are sent.

Table 31: show link-management routing Output Fields (*continued*)

Field Name	Field Description
<b>State</b>	State of the control channel.
<b>TE link name</b>	Traffic-engineered link name.
<b>State</b>	State of the traffic-engineered link: <b>Up</b> or <b>Down</b> .
<b>Local identifier</b>	Identifier of the local side of the link.
<b>Remote identifier</b>	Identifier of the remote side of the link.
<b>Local address</b>	Address of the local side of the link.
<b>Remote address</b>	Address of the remote side of the link.
<b>Encoding</b>	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b> , <b>Ethernet</b> , and <b>Packet</b> .
<b>Minimum bandwidth</b>	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
<b>Maximum bandwidth</b>	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
<b>Total bandwidth</b>	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link.
<b>Available bandwidth</b>	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
<b>Resource</b>	Forwarding adjacency LSP information.
<b>Type</b>	Type of resource. The type is always a forwarding adjacency LSP.
<b>State</b>	State of the LSP: <b>Up</b> or <b>Down</b> .
<b>System Identifier</b>	Internal identifier for the peer. The range of values is <b>0</b> through <b>64,000</b> .
<b>Total bandwidth</b>	Bandwidth resource, in bps or Mbps, on the TE-link learned from the routing process.
<b>Traffic parameters</b>	<ul style="list-style-type: none"> <li>• <b>Encoding</b>—Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b>, <b>Ethernet</b>, and <b>Packet</b>.</li> <li>• <b>Switching</b>—Type of switching that can be performed on the traffic-engineered link: <b>PSC-1</b> and <b>Packet</b>.</li> <li>• <b>Granularity</b>—Layer 2 data for switching Layer 2 LSPs for this resource. Not supported. This value is always <b>unknown</b>.</li> </ul>



## Sample Output

### show link-management routing

```

user@host> show link-management routing
Peer name: __rpd:fe-0/1/0.0, System identifier: 2147483649
State: Up, Control address: (null)
Control-channel      State
fe-0/1/0.0           Active

Peer name: __rpd:fe-0/1/2.0, System identifier: 2147483650
State: Up, Control address: (null)
Control-channel      State
fe-0/1/2.0           Active

Peer name: __rpd:so-0/2/0.0, System identifier: 2147483651
State: Down, Control address: (null)
Control-channel      State
so-0/2/0.0           State

Peer name: __rpd:so-0/2/1.0, System identifier: 2147483652
State: Down, Control address: (null)
Control-channel      State
so-0/2/1.0           State

...

TE link name: __rpd:fe-0/1/0.0, State: Up
Local identifier: 2147483649, Remote identifier: 0,
Local address: 192.168.37.66, Remote address: 192.168.37.66,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:fe-0/1/2.0, State: Up
Local identifier: 2147483650, Remote identifier: 0,
Local address: 192.168.37.73, Remote address: 192.168.37.73,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

TE link name: __rpd:so-0/2/0.0, State: Down
Local identifier: 2147483651, Remote identifier: 0,
Local address: 192.168.37.82, Remote address: 192.168.37.95,
Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 155.52Mbps,
Total bandwidth: 155.52Mbps, Available bandwidth: 155.52Mbps

...

Resource: falsp-bd, Type: LSP, State: Dn System identifier: 2147483652,
Total bandwidth: 0bps, Traffic parameters: Encoding: Packet, Switching: Packet,
Granularity: Unknown

Resource: falsp-be, Type: LSP, State: Up System identifier: 2147483654,
Total bandwidth: bw[1]=10Mbps, Traffic parameters: Encoding: Packet,
Switching: Packet, Granularity: Unknown

```

## show link-management statistics

<b>Syntax</b>	show link-management statistics <peer <name <i>name</i> >>
<b>Release Information</b>	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display statistical information for Link Management Protocol (LMP) packets.
<b>Options</b>	<b>none</b> —Display information for all peers.  <b>peer &lt;name <i>name</i>&gt;</b> —(Optional) Display information for all peers or for the specified peer only.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 337</a></li> <li>• <a href="#">show link-management peer on page 341</a></li> <li>• <a href="#">show link-management routing on page 343</a></li> <li>• <a href="#">show link-management te-link on page 348</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show link-management statistics on page 347</a>
<b>Output Fields</b>	<a href="#">Table 32 on page 346</a> describes the output fields for the <b>show link-management statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 32: show link-management statistics Output Fields**

Field Name	Field Description
Received packets	Number of received packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Received bad packets	Number of received bad packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Small packets	Number of packets that are too small.
Wrong protocol version	Number of packets specifying the wrong LMP version.
Messages for unknown peer	Number of packets destined for an unknown peer.
Messages for bad state	Number of packets indicating a state that does not match the recipient.
Stale acknowledgments	Number of <b>configAck</b> and <b>LinkSummaryAck</b> packets received that have a stale message ID.

Table 32: show link-management statistics Output Fields (*continued*)

Field Name	Field Description
<b>Stale negative acknowledgments</b>	Number of <b>configNack</b> and <b>LinkSummaryNack</b> packets received that have a stale message ID.
<b>Sent packets</b>	Number of sent packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
<b>Retransmitted packets</b>	Number of retransmitted packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
<b>Dropped packets</b>	Number of packets sent, by message type, that have been dropped by the receiver after the LMP retransmission interval has been exceeded. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.

## Sample Output

### show link-management statistics

```

user@host> show link-management statistics peer pro4-a
Statistics for peer pro4-a
  Received packets
    Config: 1
    Hello: 2572
  Small packets: 0
  Wrong protocol version: 0
  Messages for unknown peer: 0
  Messages for bad state: 0
  Stale acknowledgments: 0
  Stale negative acknowledgments: 0
  Sent packets
    Config: 2
    ConfigAck: 1
    Hello: 2572
  Retransmitted packets
    Config: 1

```

## show link-management te-link

<b>Syntax</b>	show link-management te-link <brief   detail> <name <i>name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.
<b>Options</b>	<b>none</b> —Display information for all traffic-engineered links.  <b>brief   detail</b> —(Optional) Display the specified level of output.  <b>name <i>name</i></b> —(Optional) Display information for the specified traffic-engineered link only.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 337</a></li> <li>• <a href="#">show link-management peer on page 341</a></li> <li>• <a href="#">show link-management routing on page 343</a></li> <li>• <a href="#">show link-management statistics on page 346</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show link-management te-link on page 349</a>
<b>Output Fields</b>	Table 33 on page 348 describes the output fields for the <b>show link-management te-link</b> command. Output fields are listed in the approximate order in which they appear.

**Table 33: show link-management te-link Output Fields**

Field Name	Field Description
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: <b>Up</b> or <b>Down</b> .
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b> , <b>Ethernet</b> , <b>Packet</b> , and <b>PDH</b> .

Table 33: show link-management te-link Output Fields (*continued*)

Field Name	Field Description
<b>Switching</b>	Type of switching that can be performed on the traffic-engineered link. Supported values are <b>PSC-1</b> and <b>Packet</b> .
<b>Minimum bandwidth</b>	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
<b>Maximum bandwidth</b>	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link.
<b>Total bandwidth</b>	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link (in bps).
<b>Available Bandwidth</b>	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
<b>Name</b>	Name of the interface.
<b>State</b>	State of the interface: <b>Up</b> or <b>Down</b> .
<b>Local ID</b>	Identifier of the local side of the interface.
<b>Remote ID</b>	Identifier of the remote side of the interface.
<b>Bandwidth</b>	Bandwidth, in bps or Mbps, of the member interface.
<b>Used</b>	Whether the resource is allocated to an LSP: <b>Yes</b> or <b>No</b> .
<b>LSP-name</b>	LSP name.

## Sample Output

### show link-management te-link

```

user@host> show link-management te-link
TE link name: FA-bd, State: Up
  Local identifier: 4144, Remote identifier: 0, Local address: 2.2.2.1,
  Remote address: 2.2.2.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 0bps, Total bandwidth: 0bps,
  Available bandwidth: 0bps
    Name      State Local ID Remote ID      Bandwidth Used  LSP-name
    falsp-bd  Dn      43077      0             0bps No
TE link name: FA-be, State: Up
  Local identifier: 4145, Remote identifier: 0, Local address: 1.1.1.1,
  Remote address: 1.1.1.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 10Mbps, Total bandwidth: 10Mbps,
  Available bandwidth: 8Mbps
    Name      State Local ID Remote ID      Bandwidth Used  LSP-name
    falsp-be  Up      43076      0          10Mbps Yes  e2elasp-bf

```

## show mpls call-admission-control

<b>List of Syntax</b>	<a href="#">Syntax on page 350</a> <a href="#">Syntax (EX Series Switches) on page 350</a>
<b>Syntax</b>	<pre>show mpls call-admission-control &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;lsp-name&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show mpls call-admission-control &lt;lsp-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information.
<b>Options</b>	<p><b>none</b>—Display CAC information for all LSPs.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display MPLS LSP CAC information for the specified instance. If <b>instance-name</b> is omitted, MPLS LSP CAC information for the master instance is displayed.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>lsp-name</b>—(Optional) Display CAC information for the specified LSP only.</p>
<b>Additional Information</b>	The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by any Layer 2 connection at that class type.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls call-admission-control on page 351</a>
<b>Output Fields</b>	<a href="#">Table 34 on page 350</a> describes the output fields for the <b>show mpls call-admission-control</b> command. Output fields are listed in the approximate order in which they appear.

**Table 34: show mpls call-admission-control Output Fields**

Field Name	Field Description
<b>Available bandwidth</b>	Current available bandwidth on each LSP path. Depending on whether the LSP is an E-LSP or a regular LSP, either per-class bandwidth or a single bandwidth value (corresponding to best-effort bandwidth at <b>ct0</b> ) is displayed. The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by some Layer 2 connections at that class type.

Table 34: show mpls call-admission-control Output Fields (*continued*)

Field Name	Field Description
Layer2 connections	Different Layer 2 connections that had some bandwidth requirement and were admitted into an LSP path.
LSP name	LSP pathname.
Neighbor address	Neighbor address from which CAC and bandwidth booking are configured for Layer 2 circuits.
Circuit	Interface name and circuit information.
Primary	LSP's primary standby path.
Standby	LSP's secondary standby path.
VC bandwidth	Bandwidth constraints associated with a Layer 2 circuit route.

## Sample Output

### show mpls call-admission-control

```
user@host# show mpls call-admission-control
```

```

LSP name: pro1-be
*Primary
  Available bandwidth: 0bps

LSP name: pro1-be-1
*Primary
  Available bandwidth: 60kbps

LSP name: pro1-be-gold
*Primary
  Available bandwidth: <ct0 50kbps> <ct1 20kbps> <ct2 30kbps> <ct3 0bps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.0(vc 5)
    VC bandwidth: <ct0 50kbps> <ct1 40kbps> <ct2 40kbps>

LSP name: pro1-be-gold-2
*Primary
  Available bandwidth: <ct0 0bps> <ct1 40kbps> <ct2 40kbps> <ct3 0bps>

LSP name: pro1-be-silver
*Primary  prim1
  Available bandwidth: <ct0 10kbps> <ct1 20kbps> <ct2 0bps> <ct3 40kbps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>
  Standby  sec1
  Available bandwidth: <ct0 10kbps> <ct1 10kbps> <ct2 20kbps> <ct3 0bps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>

```

## show mpls cspf

<b>List of Syntax</b>	<a href="#">Syntax on page 352</a> <a href="#">Syntax (EX Series Switches) on page 352</a>
<b>Syntax</b>	<pre>show mpls cspf &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	show mpls cspf
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics.
<b>Options</b>	<p><b>none</b>—Display MPLS CSFP statistics.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display MPLS CSPF information for the specified instance. If <i>instance-name</i> is omitted, MPLS CSPF information for the master instance is displayed.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls cspf on page 353</a>
<b>Output Fields</b>	<a href="#">Table 35 on page 352</a> describes the output fields for the <b>show mpls cspf</b> command. Output fields are listed in the approximate order in which they appear.

**Table 35: show mpls cspf Output Fields**

Field Name	Field Description
<b>Queue length</b>	Number of LSPs queued for automatic path computation.
<b>current</b>	Current queue length.
<b>maximum</b>	Maximum queue length (high-water mark).
<b>dequeued</b>	Number of aborted computation attempts.
<b>Paths</b>	Counters for label-switched path computations.
<b>total</b>	Sum of the next four fields.



Table 35: show mpls cspf Output Fields (*continued*)

Field Name	Field Description
<b>successful</b>	Number of path computations that were successfully completed.
<b>no route</b>	Number of path computations that failed because the destination is unreachable.
<b>Sys Error</b>	Number of path computations that failed because of lack of memory.
<b>CSPFs</b>	Total number of CSPF computations. A single path might require multiple CSPF computations.
<b>Time</b>	Time, in seconds, required to perform the label-switched path computation.
<b>Total</b>	Total amount of time consumed by the CSPF path computation algorithm.
<b>CSPFs</b>	Total number of CSPF computations.
<b>Avg per CSPF</b>	Average amount of time required for each CSPF computation.
<b>% of rpd</b>	Percentage of routing process CPU used in the CSPF computation.

## Sample Output

### show mpls cspf

```

user@host> show mpls cspf
CSPF statistics
Queue length  current      maximum      dequeued
              0            0            0
Paths          total      successful      no route      sys error      CSPFs
              0            0            0            0            0
Time (secs)    total      CSPFs      avg per CSPF      % of rpd
              0.000000    0.000000    0.000000          0.0000

```

## show mpls diffserv-te

<b>List of Syntax</b>	<a href="#">Syntax on page 354</a> <a href="#">Syntax (EX Series Switches) on page 354</a>
<b>Syntax</b>	<pre>show mpls diffserve-te &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	show mpls diffserve-te
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) Differentiated Services (DiffServ) class and preemption priority information.
<b>Options</b>	<p><b>none</b>—Display DiffServ classes and priorities used by MPLS LSPs.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display DiffServ classes and priorities used by MPLS LSPs for the specified instance. If <b><i>instance-name</i></b> is omitted, DiffServ information for the master instance is displayed.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls diffserv-te on page 355</a>
<b>Output Fields</b>	<p><a href="#">Table 36 on page 354</a> describes the output fields for the <b>show mpls diffserv-te</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 36: show mpls diffserv-te Output Fields**

Field Name	Field Description
<b>Bandwidth model</b>	Bandwidth constraint model supported. The maximum allocation model (MAM) for EXP-inferred LSPs (E-LSPs) is currently supported.
<b>TE class</b>	DiffServ traffic engineering class.
<b>Traffic class</b>	<p>MPLS class type that corresponds to the DiffServ traffic engineering class:</p> <ul style="list-style-type: none"> <li>• <b>ct0</b>—Best effort</li> <li>• <b>ct1</b>—Assured forwarding</li> <li>• <b>ct2</b>—Expedited forwarding</li> <li>• <b>ct3</b>—Network control</li> </ul>

Table 36: show mpls diffserv-te Output Fields (*continued*)

Field Name	Field Description
Priority	MPLS preemption priority for this class type, a value from 0 through 7. Interior gateway protocols (IGPs) distribute information about the available bandwidth for each traffic engineering class.

## Sample Output

### show mpls diffserv-te

```
user@host> show mpls diffserv-te
Bandwidth model: Maximum Allocation Model with support for E-LSPs.
TE class      Traffic class  Priority
te0           ct0           3
te1           ct1           2
```

## show route forwarding-table

<b>Syntax</b>	<pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;ccc ccc-interface-name&gt; &lt;destination&gt; &lt;family family-name&gt; &lt;label label&gt; &lt;matching ip_prefix&gt; &lt;multicast&gt; &lt;vpn vpn&gt;</pre>
<b>Release Information</b>	Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.
<b>Options</b>	<p><b>none</b>—Display the routes in the forwarding table.</p> <p><b>detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>ccc</b>—(Optional) Display the specified circuit cross-connect interface name for entries to match.</p> <p><b>destination</b>—(Optional) Display the destination prefix.</p> <p><b>family family-name</b>—(Optional) Display routing table entries for the specified family: <b>ethernet-switching</b>, <b>inet</b>, <b>inet6</b>, <b>iso</b>, <b>mpls</b>, <b>vlan classification</b>.</p> <p><b>label label</b>—(Optional) Display route entries for the specified label name.</p> <p><b>matching ip_prefix</b>—(Optional) Display route entries for the specified IP prefix.</p> <p><b>multicast</b>—(Optional) Display route entries for multicast routes.</p> <p><b>vpn vpn</b>—(Optional) Display route entries for the specified VPN.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring MPLS on EX8200 and EX4500 Switches</i></li> <li>• <i>Configuring MPLS on EX8200 and EX4500 Provider Switches (CLI Procedure)</i></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show route forwarding-table on page 358</a></p> <p><a href="#">show route forwarding-table summary on page 359</a></p> <p><a href="#">show route forwarding-table extensive on page 359</a></p> <p><a href="#">show route forwarding-table ccc on page 361</a></p> <p><a href="#">show route forwarding-table family (MPLS) on page 361</a></p>

[show route forwarding-table family \(IPv6\) on page 361](#)

[show route forwarding-table label on page 362](#)

[show route forwarding-table matching on page 362](#)

[show route forwarding-table multicast on page 362](#)

**Output Fields** Table 37 on page 357 lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified or when the **detail** keyword is used instead of the **extensive** keyword.

**Table 37: show route forwarding-table Output Fields**

Field Name	Field Description	Level of Output
<b>Routing table</b>	Name of the routing table (for example, <b>inet</b> , <b>inet6</b> , <b>mpls</b> ).	All levels
<b>Address family</b>	Address family (for example, <b>IP</b> , <b>IPv6</b> , <b>ISO</b> , <b>MPLS</b> ).	All levels
<b>Destination</b>	Destination of the route.	<b>detail</b> , <b>extensive</b>
<b>Route Type (Type)</b>	How the route was placed into the forwarding table. When the <b>detail</b> keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> <li>• <b>cloned (clon)</b>—(TCP or multicast only) Cloned route.</li> <li>• <b>destination (dest)</b>—Remote addresses directly reachable through an interface.</li> <li>• <b>destination down (iddn)</b>—Destination route for which the interface is unreachable.</li> <li>• <b>interface cloned (ifcl)</b>—Cloned route for which the interface is unreachable.</li> <li>• <b>route down (ifdn)</b>—Interface route for which the interface is unreachable.</li> <li>• <b>ignore (ignr)</b>—Ignore this route.</li> <li>• <b>interface (intf)</b>—Installed as a result of configuring an interface.</li> <li>• <b>permanent (perm)</b>—Routes installed by the kernel when the routing table is initialized.</li> <li>• <b>user</b>—Routes installed by the routing protocol process or as a result of the configuration.</li> </ul>	All levels
<b>Route reference (RtRef)</b>	Number of routes to reference.	<b>detail</b> , <b>extensive</b>
<b>Flags</b>	Route type flags: <ul style="list-style-type: none"> <li>• <b>none</b>—No flags are enabled.</li> <li>• <b>accounting</b>—Route has accounting enabled.</li> <li>• <b>cached</b>—Cache route.</li> <li>• <b>incoming-iface interface-number</b>—Check against incoming interface.</li> <li>• <b>prefix load balance</b>—Load balancing is enabled for this prefix.</li> <li>• <b>sent to PFE</b>—Route has been sent to the Packet Forwarding Engine.</li> <li>• <b>static</b>—Static route.</li> </ul>	<b>extensive</b>
<b>Nexthop</b>	IP address of the next hop to the destination.	<b>detail</b> , <b>extensive</b>

Table 37: show route forwarding-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Next hop type (Type)</b>	<p>Next-hop type. When the <b>detail</b> keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> <li>• <b>broadcast (bcst)</b>—Broadcast.</li> <li>• <b>deny</b>—Deny.</li> <li>• <b>hold</b>—Next hop is waiting to be resolved into a unicast or multicast type.</li> <li>• <b>indexed (idxd)</b>—Indexed next hop.</li> <li>• <b>indirect (indr)</b>—Indirect next hop.</li> <li>• <b>local (locl)</b>—Local address on an interface.</li> <li>• <b>routed multicast (mcrst)</b>—Regular multicast next hop</li> <li>• <b>multicast (mcst)</b>—Wire multicast next hop (limited to the LAN).</li> <li>• <b>multicast discard (mdsc)</b>—Multicast discard.</li> <li>• <b>multicast group (mgrp)</b> —Multicast group member.</li> <li>• <b>receive (recv)</b>—Receive.</li> <li>• <b>reject (rjct)</b>—Discard. An ICMP unreachable message was sent.</li> <li>• <b>resolve (rslv)</b>—Resolving the next hop.</li> <li>• <b>unicast (ucst)</b>—Unicast.</li> <li>• <b>unilist (ulst)</b>—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.</li> </ul>	<b>detail, extensive</b>
<b>Index</b>	Software index of the next hop that is used to route the traffic for a given prefix.	<b>detail, extensive none</b>
<b>Route interface-index</b>	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	<b>extensive</b>
<b>Reference (NhRef)</b>	Number of routes that refer to this next hop.	<b>none detail, extensive</b>
<b>Next-hop interface (Netif)</b>	Interface used to reach the next hop.	<b>none detail, extensive</b>
<b>Alternate forward nh index</b>	Index number of the alternate next hop interface. Seen with <b>multicast</b> option only.	<b>extensive</b>
<b>Next-hop L3 Interface</b>	The next hop layer 3 interface. This option can be expressed as a VLAN name and is only seen with the <b>multicast</b> option.	<b>extensive</b>
<b>Next-hop L2 Interfaces</b>	The next hop layer 2 interfaces. Seen with <b>multicast</b> option only.	<b>extensive</b>

## Sample Output

### show route forwarding-table

```

user@switch> show route forwarding-table

Routing table: default.inet

```

```

Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          user   2 0:12:f2:21:cf:0    ucst  333   5 me0.0
default          perm   0                               rjct   36   2
0.0.0.0/32       perm   0                               dscd   34   1
2.2.2.0/24       intf   0                               rslv  1309   1 ae0.0
2.2.2.0/32       dest   0 2.2.2.0           recv  1307   1 ae0.0
2.2.2.1/32       dest   0 0:21:59:cc:89:c0  ucst  1320   1 ae0.0
2.2.2.2/32       intf   0 2.2.2.2           locl  1308   2
2.2.2.2/32       dest   0 2.2.2.2           locl  1308   2
2.2.2.255/32     dest   0 2.2.2.255         bcst  1306   1 ae0.0
3.3.3.0/24       intf   0                               rslv  1313   1 ae1.0
3.3.3.0/32       dest   0 3.3.3.0           recv  1311   1 ae1.0
3.3.3.1/32       intf   0 3.3.3.1           locl  1312   2
3.3.3.1/32       dest   0 3.3.3.1           locl  1312   2
3.3.3.2/32       dest   0 0:21:59:cc:89:c1  ucst  1321  24 ae1.0
3.3.3.255/32     dest   0 3.3.3.255         bcst  1310   1 ae1.0
4.4.4.0/24       user   0 3.3.3.2           ucst  1321  24 ae1.0
8.8.8.8/32       user   0 3.3.3.2           ucst  1321  24 ae1.0
9.9.9.9/32       intf   0 9.9.9.9           locl  1280   1
10.10.10.10/32   user   0 3.3.3.2           ucst  1321  24 ae1.0
10.93.8.0/21     intf   0                               rslv  323   1 me0.0
10.93.8.0/32     dest   0 10.93.8.0         recv  321   1 me0.0
10.93.13.238/32  intf   0 10.93.13.238      locl  322   2
10.93.13.238/32  dest   0 10.93.13.238      locl  322   2
10.93.15.254/32  dest   0 0:12:f2:21:cf:0    ucst  333   5 me0.0
10.93.15.255/32  dest   0 10.93.15.255      bcst  320   1 me0.0
14.14.14.0/24    ifdn   0                               rslv  1319   1 ge-0/0/25.0
14.14.14.0/32    iddn   0 14.14.14.0        recv  1317   1 ge-0/0/25.0
14.14.14.2/32    user   0                               rjct   36   2
14.14.14.2/32    intf   0 14.14.14.2        locl  1318   2
14.14.14.2/32    iddn   0 14.14.14.2        locl  1318   2
14.14.14.255/32  iddn   0 14.14.14.255      bcst  1316   1 ge-0/0/25.0
224.0.0.0/4      perm   1                               mdsc   35   1
224.0.0.1/32     perm   0 224.0.0.1         mcst   31   3
224.0.0.5/32     user   1 224.0.0.5         mcst   31   3
255.255.255.255/32 perm   0                               bcst   32   1

```

### show route forwarding-table summary

```
user@switch> show route forwarding-table summary
```

```

Routing table: default.inet
Internet:
    user:          6 routes
    perm:          5 routes
    intf:          8 routes
    dest:         12 routes
    ifdn:          1 routes
    iddn:          3 routes

```

### show route forwarding-table extensive

```
user@switch> show route forwarding-table summary
```

```

Routing table: default.inet [Index 0]
Internet:

Destination: default
Route type: user
Route reference: 2
Route interface-index: 0

```

```

Flags: sent to PFE, rt nh decoupled
Nexthop: 0:12:f2:21:cf:0
Next-hop type: unicast          Index: 333      Reference: 5
Next-hop interface: me0.0

Destination: default
Route type: permanent
Route reference: 0              Route interface-index: 0
Flags: none
Next-hop type: reject          Index: 36      Reference: 2

Destination: 0.0.0.0/32
Route type: permanent
Route reference: 0              Route interface-index: 0
Flags: sent to PFE
Next-hop type: discard         Index: 34      Reference: 1

Destination: 2.2.2.0/24
Route type: interface
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Next-hop type: resolve         Index: 1309    Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.0/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Nexthop: 2.2.2.0
Next-hop type: receive         Index: 1307    Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.1/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Nexthop: 0:21:59:cc:89:c0
Next-hop type: unicast         Index: 1320    Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.2/32
Route type: interface
Route reference: 0              Route interface-index: 0
Flags: sent to PFE
Nexthop: 2.2.2.2
Next-hop type: local           Index: 1308    Reference: 2

Destination: 2.2.2.2/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: none
Nexthop: 2.2.2.2
Next-hop type: local           Index: 1308    Reference: 2

Destination: 2.2.2.255/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Nexthop: 2.2.2.255
Next-hop type: broadcast       Index: 1306    Reference: 1
Next-hop interface: ae0.0

```



**show route forwarding-table ccc**

```

user@switch> show route forwarding-table ccc ge-0/0/0.10
Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
ge-0/0/0.10      (CCC) user    0 3.3.3.2      Push 300112 1343  2 ae1.0

```

**show route forwarding-table family (MPLS)**

```

user@switch> show route forwarding-table family mpls

Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm    0
0                user    0      recv    49    3
1                user    0      recv    49    3
2                user    0      recv    49    3
299776           user    0      Pop     1334   2 ge-0/0/0.10
299792           user    0      Pop     1339   2 ge-0/0/0.14
299808           user    0      Pop     1341   2 ge-0/0/0.2
299824           user    0      Pop     1344   2 ge-0/0/0.11
299840           user    0      Pop     1345   2 ge-0/0/0.13
299856           user    0      Pop     1346   2 ge-0/0/0.18
299872           user    0      Pop     1347   2 ge-0/0/0.16
299888           user    0      Pop     1348   2 ge-0/0/0.7
299904           user    0      Pop     1349   2 ge-0/0/0.20
299920           user    0      Pop     1350   2 ge-0/0/0.19
299936           user    0      Pop     1351   2 ge-0/0/0.17
299952           user    0      Pop     1352   2 ge-0/0/0.9
299968           user    0      Pop     1353   2 ge-0/0/0.1
299984           user    0      Pop     1354   2 ge-0/0/0.12
300000           user    0      Pop     1355   2 ge-0/0/0.8
300016           user    0      Pop     1356   2 ge-0/0/0.4
300032           user    0      Pop     1357   2 ge-0/0/0.5
300048           user    0      Pop     1358   2 ge-0/0/0.3
300064           user    0      Pop     1359   2 ge-0/0/0.15
ge-0/0/0.1       (CCC) user    0 3.3.3.2      Push 300064 1340  2 ae1.0
ge-0/0/0.2       (CCC) user    0 3.3.3.2      Push 299872 1328  2 ae1.0
ge-0/0/0.3       (CCC) user    0 3.3.3.2      Push 299792 1323  2 ae1.0
ge-0/0/0.4       (CCC) user    0 3.3.3.2      Push 300016 1337  2 ae1.0
ge-0/0/0.5       (CCC) user    0 3.3.3.2      Push 299824 1325  2 ae1.0
ge-0/0/0.7       (CCC) user    0 3.3.3.2      Push 299920 1331  2 ae1.0
ge-0/0/0.8       (CCC) user    0 3.3.3.2      Push 299840 1326  2 ae1.0
ge-0/0/0.9       (CCC) user    0 3.3.3.2      Push 299888 1329  2 ae1.0
ge-0/0/0.10      (CCC) user    0 3.3.3.2      Push 300112 1343  2 ae1.0
ge-0/0/0.11      (CCC) user    0 3.3.3.2      Push 299776 1322  2 ae1.0
ge-0/0/0.12      (CCC) user    0 3.3.3.2      Push 299952 1333  2 ae1.0
ge-0/0/0.13      (CCC) user    0 3.3.3.2      Push 300096 1342  2 ae1.0
ge-0/0/0.14      (CCC) user    0 3.3.3.2      Push 299984 1335  2 ae1.0
ge-0/0/0.15      (CCC) user    0 3.3.3.2      Push 299936 1332  2 ae1.0
ge-0/0/0.16      (CCC) user    0 3.3.3.2      Push 299808 1324  2 ae1.0
ge-0/0/0.17      (CCC) user    0 3.3.3.2      Push 300000 1336  2 ae1.0
ge-0/0/0.18      (CCC) user    0 3.3.3.2      Push 300032 1338  2 ae1.0
ge-0/0/0.19      (CCC) user    0 3.3.3.2      Push 299904 1330  2 ae1.0
ge-0/0/0.20      (CCC) user    0 3.3.3.2      Push 299856 1327  2 ae1.0

```

**show route forwarding-table family (IPv6)**

```

user@switch> show route forwarding-table family inet6

```

```

Routing table: default.inet6
Internet6:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0                rjct  44    1
::/128           perm  0                dscd  42    1
ff00::/8         perm  0                mdsc  43    1
ff02::1/128     perm  0 ff02::1       mcst  39    1

```

```

Routing table: default-switch.inet6
Internet6:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0                rjct  530   1
::/128           perm  0                dscd  528   1
2:1::3a00/312   user  0                indr 131070 2
                comp  572   1
2:1::3a82/320   user  0                indr 131071 3
                comp  573   1
2:1::3af0/320   user  0                indr 131071 3
                comp  573   1
2:1:0:ff00::/56 user  0                mdsc  529   2
ff00::/8         perm  0                mdsc  529   2
ff02::1/128     perm  0 ff02::1       mcst  526   1

```

```

Routing table: __master.anon__.inet6
Internet6:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0                rjct  554   1
::/128           perm  0                dscd  552   1
ff00::/8         perm  0                mdsc  553   1
ff02::1/128     perm  0 ff02::1       mcst  550   1

```

### show route forwarding-table label

```
user@switch> show route forwarding-table label 29976
```

```

Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
299776           user  0                Pop   1334   2 ge-0/0/0.10

```

### show route forwarding-table matching

```
user@switch> show route forwarding-table matching 3
```

```

Routing table: default.inet
Internet:

```

### show route forwarding-table multicast

```
user@switch> show route forwarding-table multicast
```

```

Routing table: default.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
224.0.0.0/4       perm  1                mdsc  35    1
224.0.0.1/32      perm  0 224.0.0.1       mcst  31    3
224.0.0.5/32      user  1 224.0.0.5       mcst  31    3

```

```

Routing table: __master.anon__.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
224.0.0.0/4       perm  0                mdsc 1289   1

```

```
224.0.0.1/32      perm      0 224.0.0.1      mcst  1285      1
```

```
Routing table: default.inet6
```

```
Internet6:
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
ff00::/8	perm	0		mdsc	43	1	
ff02::1/128	perm	0	ff02::1	mcst	39	1	

## show mpls interface

<b>List of Syntax</b>	<a href="#">Syntax on page 364</a> <a href="#">Syntax (EX Series Switches) on page 364</a>
<b>Syntax</b>	<pre>show mpls interface &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	show mpls interface
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces.
<b>Options</b>	<p><b>none</b>—Display information about MPLS-enabled interfaces.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about MPLS-enabled interfaces for the specified routing instance. If <b>instance-name</b> is omitted, information about MPLS-enabled interfaces is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Additional Information</b>	MPLS is enabled on an interface when the interface is configured with both the <b>set protocol mpls interface <i>interface-name</i></b> and <b>set interface <i>interface-name</i> unit 0 family mpls</b> statements.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls interface on page 365</a>
<b>Output Fields</b>	<p><a href="#">Table 38 on page 364</a> describes the output fields for the <b>show mpls interface</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 38: show mpls interface Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of the interface.
<b>State</b>	State of the interface: <b>Up</b> or <b>Dn</b> (down).
<b>Administrative groups</b>	Administratively assigned colors of the link.

Table 38: show mpls interface Output Fields (*continued*)

Field Name	Field Description
Maximum labels	Maximum number of MPLS labels upon which MPLS can operate on a logical interface. This is configured using the <b>maximum-labels</b> statement at the [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] or the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] hierarchy levels.
Static protection revert time	Time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path. This is configured using the <b>protection-revert-time</b> statement at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels.
Always mark connection protection tlv	Enabled or Disabled: Enabled indicates that the <b>always-mark-connection-protection-tlv</b> statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. When this statement is configured, it marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, the <b>switch-away-lsps</b> statement must be configured.
Switch away lsps	Enabled or Disabled: Enabled indicates that the <b>switch-away-lsps</b> statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. This enables you to switch an LSP away from a network node using a bypass LSP. This feature can be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.

## Sample Output

### show mpls interface

```
user@host> show mpls interface
```

```
Interface: ge-0/2/1.57
State: Up
Administrative group: <none>
Maximum labels: 5
Static protection revert time: 5 seconds
Always mark connection protection tlv: Disabled
Switch away lsps : Disabled
```

## show link-management statistics

<b>Syntax</b>	show link-management statistics <peer <name <i>name</i> >>
<b>Release Information</b>	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display statistical information for Link Management Protocol (LMP) packets.
<b>Options</b>	<b>none</b> —Display information for all peers.  <b>peer &lt;name <i>name</i>&gt;</b> —(Optional) Display information for all peers or for the specified peer only.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 337</a></li> <li>• <a href="#">show link-management peer on page 341</a></li> <li>• <a href="#">show link-management routing on page 343</a></li> <li>• <a href="#">show link-management te-link on page 348</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show link-management statistics on page 367</a>
<b>Output Fields</b>	<a href="#">Table 32 on page 346</a> describes the output fields for the <b>show link-management statistics</b> command. Output fields are listed in the approximate order in which they appear.

**Table 39: show link-management statistics Output Fields**

Field Name	Field Description
Received packets	Number of received packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Received bad packets	Number of received bad packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Small packets	Number of packets that are too small.
Wrong protocol version	Number of packets specifying the wrong LMP version.
Messages for unknown peer	Number of packets destined for an unknown peer.
Messages for bad state	Number of packets indicating a state that does not match the recipient.
Stale acknowledgments	Number of <b>configAck</b> and <b>LinkSummaryAck</b> packets received that have a stale message ID.

Table 39: show link-management statistics Output Fields (*continued*)

Field Name	Field Description
<b>Stale negative acknowledgments</b>	Number of <b>configNack</b> and <b>LinkSummaryNack</b> packets received that have a stale message ID.
<b>Sent packets</b>	Number of sent packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
<b>Retransmitted packets</b>	Number of retransmitted packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
<b>Dropped packets</b>	Number of packets sent, by message type, that have been dropped by the receiver after the LMP retransmission interval has been exceeded. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.

## Sample Output

### show link-management statistics

```

user@host> show link-management statistics peer pro4-a
Statistics for peer pro4-a
  Received packets
    Config: 1
    Hello: 2572
  Small packets: 0
  Wrong protocol version: 0
  Messages for unknown peer: 0
  Messages for bad state: 0
  Stale acknowledgments: 0
  Stale negative acknowledgments: 0
  Sent packets
    Config: 2
    ConfigAck: 1
    Hello: 2572
  Retransmitted packets
    Config: 1

```

## show link-management te-link

<b>Syntax</b>	show link-management te-link <brief   detail> <name <i>name</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.
<b>Options</b>	<b>none</b> —Display information for all traffic-engineered links.  <b>brief   detail</b> —(Optional) Display the specified level of output.  <b>name <i>name</i></b> —(Optional) Display information for the specified traffic-engineered link only.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show link-management on page 337</a></li> <li>• <a href="#">show link-management peer on page 341</a></li> <li>• <a href="#">show link-management routing on page 343</a></li> <li>• <a href="#">show link-management statistics on page 346</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show link-management te-link on page 369</a>
<b>Output Fields</b>	Table 33 on page 348 describes the output fields for the <b>show link-management te-link</b> command. Output fields are listed in the approximate order in which they appear.

Table 40: show link-management te-link Output Fields

Field Name	Field Description
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: <b>Up</b> or <b>Down</b> .
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include <b>SDH/SONET</b> , <b>Ethernet</b> , <b>Packet</b> , and <b>PDH</b> .



Table 40: show link-management te-link Output Fields (*continued*)

Field Name	Field Description
<b>Switching</b>	Type of switching that can be performed on the traffic-engineered link. Supported values are <b>PSC-1</b> and <b>Packet</b> .
<b>Minimum bandwidth</b>	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
<b>Maximum bandwidth</b>	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link.
<b>Total bandwidth</b>	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link (in bps).
<b>Available Bandwidth</b>	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
<b>Name</b>	Name of the interface.
<b>State</b>	State of the interface: <b>Up</b> or <b>Down</b> .
<b>Local ID</b>	Identifier of the local side of the interface.
<b>Remote ID</b>	Identifier of the remote side of the interface.
<b>Bandwidth</b>	Bandwidth, in bps or Mbps, of the member interface.
<b>Used</b>	Whether the resource is allocated to an LSP: <b>Yes</b> or <b>No</b> .
<b>LSP-name</b>	LSP name.

## Sample Output

### show link-management te-link

```

user@host> show link-management te-link
TE link name: FA-bd, State: Up
  Local identifier: 4144, Remote identifier: 0, Local address: 2.2.2.1,
  Remote address: 2.2.2.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 0bps, Total bandwidth: 0bps,
  Available bandwidth: 0bps
    Name      State Local ID Remote ID      Bandwidth Used  LSP-name
    falsp-bd   Dn      43077      0           0bps No
TE link name: FA-be, State: Up
  Local identifier: 4145, Remote identifier: 0, Local address: 1.1.1.1,
  Remote address: 1.1.1.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 10Mbps, Total bandwidth: 10Mbps,
  Available bandwidth: 8Mbps
    Name      State Local ID Remote ID      Bandwidth Used  LSP-name
    falsp-be   Up      43076      0          10Mbps Yes  e2elasp-bf

```

## show mpls call-admission-control

<b>List of Syntax</b>	<a href="#">Syntax on page 370</a> <a href="#">Syntax (EX Series Switches) on page 370</a>
<b>Syntax</b>	<pre>show mpls call-admission-control &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;lsp-name&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show mpls call-admission-control &lt;lsp-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information.
<b>Options</b>	<p><b>none</b>—Display CAC information for all LSPs.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display MPLS LSP CAC information for the specified instance. If <b>instance-name</b> is omitted, MPLS LSP CAC information for the master instance is displayed.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>lsp-name</b>—(Optional) Display CAC information for the specified LSP only.</p>
<b>Additional Information</b>	The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by any Layer 2 connection at that class type.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls call-admission-control on page 371</a>
<b>Output Fields</b>	<a href="#">Table 34 on page 350</a> describes the output fields for the <b>show mpls call-admission-control</b> command. Output fields are listed in the approximate order in which they appear.

**Table 41: show mpls call-admission-control Output Fields**

Field Name	Field Description
Available bandwidth	Current available bandwidth on each LSP path. Depending on whether the LSP is an E-LSP or a regular LSP, either per-class bandwidth or a single bandwidth value (corresponding to best-effort bandwidth at <b>ct0</b> ) is displayed. The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by some Layer 2 connections at that class type.

Table 41: show mpls call-admission-control Output Fields (*continued*)

Field Name	Field Description
Layer2 connections	Different Layer 2 connections that had some bandwidth requirement and were admitted into an LSP path.
LSP name	LSP pathname.
Neighbor address	Neighbor address from which CAC and bandwidth booking are configured for Layer 2 circuits.
Circuit	Interface name and circuit information.
Primary	LSP's primary standby path.
Standby	LSP's secondary standby path.
VC bandwidth	Bandwidth constraints associated with a Layer 2 circuit route.

## Sample Output

### show mpls call-admission-control

```
user@host# show mpls call-admission-control
```

```

LSP name: pro1-be
*Primary
  Available bandwidth: 0bps

LSP name: pro1-be-1
*Primary
  Available bandwidth: 60kbps

LSP name: pro1-be-gold
*Primary
  Available bandwidth: <ct0 50kbps> <ct1 20kbps> <ct2 30kbps> <ct3 0bps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.0(vc 5)
    VC bandwidth: <ct0 50kbps> <ct1 40kbps> <ct2 40kbps>

LSP name: pro1-be-gold-2
*Primary
  Available bandwidth: <ct0 0bps> <ct1 40kbps> <ct2 40kbps> <ct3 0bps>

LSP name: pro1-be-silver
*Primary  prim1
  Available bandwidth: <ct0 10kbps> <ct1 20kbps> <ct2 0bps> <ct3 40kbps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>
  Standby  sec1
  Available bandwidth: <ct0 10kbps> <ct1 10kbps> <ct2 20kbps> <ct3 0bps>
  Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>

```

## show mpls cspf

<b>List of Syntax</b>	<a href="#">Syntax on page 372</a> <a href="#">Syntax (EX Series Switches) on page 372</a>
<b>Syntax</b>	<pre>show mpls cspf &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	show mpls cspf
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics.
<b>Options</b>	<p><b>none</b>—Display MPLS CSFP statistics.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display MPLS CSPF information for the specified instance. If <i>instance-name</i> is omitted, MPLS CSPF information for the master instance is displayed.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls cspf on page 373</a>
<b>Output Fields</b>	<a href="#">Table 35 on page 352</a> describes the output fields for the <b>show mpls cspf</b> command. Output fields are listed in the approximate order in which they appear.

**Table 42: show mpls cspf Output Fields**

Field Name	Field Description
<b>Queue length</b>	Number of LSPs queued for automatic path computation.
<b>current</b>	Current queue length.
<b>maximum</b>	Maximum queue length (high-water mark).
<b>dequeued</b>	Number of aborted computation attempts.
<b>Paths</b>	Counters for label-switched path computations.
<b>total</b>	Sum of the next four fields.

Table 42: show mpls cspf Output Fields (*continued*)

Field Name	Field Description
<b>successful</b>	Number of path computations that were successfully completed.
<b>no route</b>	Number of path computations that failed because the destination is unreachable.
<b>Sys Error</b>	Number of path computations that failed because of lack of memory.
<b>CSPFs</b>	Total number of CSPF computations. A single path might require multiple CSPF computations.
<b>Time</b>	Time, in seconds, required to perform the label-switched path computation.
<b>Total</b>	Total amount of time consumed by the CSPF path computation algorithm.
<b>CSPFs</b>	Total number of CSPF computations.
<b>Avg per CSPF</b>	Average amount of time required for each CSPF computation.
<b>% of rpd</b>	Percentage of routing process CPU used in the CSPF computation.

## Sample Output

### show mpls cspf

```

user@host> show mpls cspf
CSPF statistics
Queue length  current      maximum      dequeued
              0            0            0
Paths          total      successful      no route      sys error      CSPFs
              0            0            0            0            0
Time (secs)    total      CSPFs      avg per CSPF      % of rpd
              0.000000    0.000000    0.000000    0.0000

```

## show mpls diffserv-te

<b>List of Syntax</b>	<a href="#">Syntax on page 374</a> <a href="#">Syntax (EX Series Switches) on page 374</a>
<b>Syntax</b>	<pre>show mpls diffserve-te &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	show mpls diffserve-te
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) Differentiated Services (DiffServ) class and preemption priority information.
<b>Options</b>	<p><b>none</b>—Display DiffServ classes and priorities used by MPLS LSPs.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display DiffServ classes and priorities used by MPLS LSPs for the specified instance. If <b><i>instance-name</i></b> is omitted, DiffServ information for the master instance is displayed.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls diffserv-te on page 375</a>
<b>Output Fields</b>	<p><a href="#">Table 36 on page 354</a> describes the output fields for the <b>show mpls diffserv-te</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 43: show mpls diffserv-te Output Fields**

Field Name	Field Description
<b>Bandwidth model</b>	Bandwidth constraint model supported. The maximum allocation model (MAM) for EXP-inferred LSPs (E-LSPs) is currently supported.
<b>TE class</b>	DiffServ traffic engineering class.
<b>Traffic class</b>	<p>MPLS class type that corresponds to the DiffServ traffic engineering class:</p> <ul style="list-style-type: none"> <li>• <b>ct0</b>—Best effort</li> <li>• <b>ct1</b>—Assured forwarding</li> <li>• <b>ct2</b>—Expedited forwarding</li> <li>• <b>ct3</b>—Network control</li> </ul>

Table 43: show mpls diffserv-te Output Fields (*continued*)

Field Name	Field Description
Priority	MPLS preemption priority for this class type, a value from 0 through 7. Interior gateway protocols (IGPs) distribute information about the available bandwidth for each traffic engineering class.

## Sample Output

### show mpls diffserv-te

```
user@host> show mpls diffserv-te
Bandwidth model: Maximum Allocation Model with support for E-LSPs.
TE class      Traffic class  Priority
te0           ct0           3
te1           ct1           2
```

## show route forwarding-table

<b>Syntax</b>	<pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;ccc ccc-interface-name&gt; &lt;destination&gt; &lt;family family-name&gt; &lt;label label&gt; &lt;matching ip_prefix&gt; &lt;multicast&gt; &lt;vpn vpn&gt;</pre>
<b>Release Information</b>	Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.
<b>Options</b>	<p><b>none</b>—Display the routes in the forwarding table.</p> <p><b>detail   extensive   summary</b>—(Optional) Display the specified level of output.</p> <p><b>ccc</b>—(Optional) Display the specified circuit cross-connect interface name for entries to match.</p> <p><b>destination</b>—(Optional) Display the destination prefix.</p> <p><b>family family-name</b>—(Optional) Display routing table entries for the specified family: <b>ethernet-switching</b>, <b>inet</b>, <b>inet6</b>, <b>iso</b>, <b>mpls</b>, <b>vlan classification</b>.</p> <p><b>label label</b>—(Optional) Display route entries for the specified label name.</p> <p><b>matching ip_prefix</b>—(Optional) Display route entries for the specified IP prefix.</p> <p><b>multicast</b>—(Optional) Display route entries for multicast routes.</p> <p><b>vpn vpn</b>—(Optional) Display route entries for the specified VPN.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Example: Configuring MPLS on EX8200 and EX4500 Switches</i></li> <li>• <i>Configuring MPLS on EX8200 and EX4500 Provider Switches (CLI Procedure)</i></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show route forwarding-table on page 378</a></p> <p><a href="#">show route forwarding-table summary on page 379</a></p> <p><a href="#">show route forwarding-table extensive on page 379</a></p> <p><a href="#">show route forwarding-table ccc on page 381</a></p> <p><a href="#">show route forwarding-table family (MPLS) on page 381</a></p>



[show route forwarding-table family \(IPv6\) on page 381](#)  
[show route forwarding-table label on page 382](#)  
[show route forwarding-table matching on page 382](#)  
[show route forwarding-table multicast on page 382](#)

**Output Fields** Table 37 on page 357 lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified or when the **detail** keyword is used instead of the **extensive** keyword.

**Table 44: show route forwarding-table Output Fields**

Field Name	Field Description	Level of Output
<b>Routing table</b>	Name of the routing table (for example, <b>inet</b> , <b>inet6</b> , <b>mpls</b> ).	All levels
<b>Address family</b>	Address family (for example, <b>IP</b> , <b>IPv6</b> , <b>ISO</b> , <b>MPLS</b> ).	All levels
<b>Destination</b>	Destination of the route.	<b>detail</b> , <b>extensive</b>
<b>Route Type (Type)</b>	How the route was placed into the forwarding table. When the <b>detail</b> keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> <li>• <b>cloned (clon)</b>—(TCP or multicast only) Cloned route.</li> <li>• <b>destination (dest)</b>—Remote addresses directly reachable through an interface.</li> <li>• <b>destination down (iddn)</b>—Destination route for which the interface is unreachable.</li> <li>• <b>interface cloned (ifcl)</b>—Cloned route for which the interface is unreachable.</li> <li>• <b>route down (ifdn)</b>—Interface route for which the interface is unreachable.</li> <li>• <b>ignore (ignr)</b>—Ignore this route.</li> <li>• <b>interface (intf)</b>—Installed as a result of configuring an interface.</li> <li>• <b>permanent (perm)</b>—Routes installed by the kernel when the routing table is initialized.</li> <li>• <b>user</b>—Routes installed by the routing protocol process or as a result of the configuration.</li> </ul>	All levels
<b>Route reference (RtRef)</b>	Number of routes to reference.	<b>detail</b> , <b>extensive</b>
<b>Flags</b>	Route type flags: <ul style="list-style-type: none"> <li>• <b>none</b>—No flags are enabled.</li> <li>• <b>accounting</b>—Route has accounting enabled.</li> <li>• <b>cached</b>—Cache route.</li> <li>• <b>incoming-iface interface-number</b>—Check against incoming interface.</li> <li>• <b>prefix load balance</b>—Load balancing is enabled for this prefix.</li> <li>• <b>sent to PFE</b>—Route has been sent to the Packet Forwarding Engine.</li> <li>• <b>static</b>—Static route.</li> </ul>	<b>extensive</b>
<b>Nexthop</b>	IP address of the next hop to the destination.	<b>detail</b> , <b>extensive</b>

Table 44: show route forwarding-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Next hop type (Type)</b>	<p>Next-hop type. When the <b>detail</b> keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> <li>• <b>broadcast (bcst)</b>—Broadcast.</li> <li>• <b>deny</b>—Deny.</li> <li>• <b>hold</b>—Next hop is waiting to be resolved into a unicast or multicast type.</li> <li>• <b>indexed (idxd)</b>—Indexed next hop.</li> <li>• <b>indirect (indr)</b>—Indirect next hop.</li> <li>• <b>local (locl)</b>—Local address on an interface.</li> <li>• <b>routed multicast (mcrst)</b>—Regular multicast next hop</li> <li>• <b>multicast (mcst)</b>—Wire multicast next hop (limited to the LAN).</li> <li>• <b>multicast discard (mdsc)</b>—Multicast discard.</li> <li>• <b>multicast group (mgrp)</b> —Multicast group member.</li> <li>• <b>receive (recv)</b>—Receive.</li> <li>• <b>reject (rjct)</b>—Discard. An ICMP unreachable message was sent.</li> <li>• <b>resolve (rslv)</b>—Resolving the next hop.</li> <li>• <b>unicast (ucst)</b>—Unicast.</li> <li>• <b>unilist (ulst)</b>—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.</li> </ul>	<b>detail, extensive</b>
<b>Index</b>	Software index of the next hop that is used to route the traffic for a given prefix.	<b>detail, extensive none</b>
<b>Route interface-index</b>	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	<b>extensive</b>
<b>Reference (NhRef)</b>	Number of routes that refer to this next hop.	<b>none detail, extensive</b>
<b>Next-hop interface (Netif)</b>	Interface used to reach the next hop.	<b>none detail, extensive</b>
<b>Alternate forward nh index</b>	Index number of the alternate next hop interface. Seen with <b>multicast</b> option only.	<b>extensive</b>
<b>Next-hop L3 Interface</b>	The next hop layer 3 interface. This option can be expressed as a VLAN name and is only seen with the <b>multicast</b> option.	<b>extensive</b>
<b>Next-hop L2 Interfaces</b>	The next hop layer 2 interfaces. Seen with <b>multicast</b> option only.	<b>extensive</b>

## Sample Output

### show route forwarding-table

```

user@switch> show route forwarding-table

Routing table: default.inet

```

```

Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          user   2 0:12:f2:21:cf:0    ucst  333   5 me0.0
default          perm   0                               rjct   36   2
0.0.0.0/32       perm   0                               dscd   34   1
2.2.2.0/24       intf   0                               rslv  1309  1 ae0.0
2.2.2.0/32       dest   0 2.2.2.0           recv  1307  1 ae0.0
2.2.2.1/32       dest   0 0:21:59:cc:89:c0  ucst  1320  1 ae0.0
2.2.2.2/32       intf   0 2.2.2.2           locl  1308  2
2.2.2.2/32       dest   0 2.2.2.2           locl  1308  2
2.2.2.255/32     dest   0 2.2.2.255        bcst  1306  1 ae0.0
3.3.3.0/24       intf   0                               rslv  1313  1 ae1.0
3.3.3.0/32       dest   0 3.3.3.0           recv  1311  1 ae1.0
3.3.3.1/32       intf   0 3.3.3.1           locl  1312  2
3.3.3.1/32       dest   0 3.3.3.1           locl  1312  2
3.3.3.2/32       dest   0 0:21:59:cc:89:c1  ucst  1321  24 ae1.0
3.3.3.255/32     dest   0 3.3.3.255        bcst  1310  1 ae1.0
4.4.4.0/24       user   0 3.3.3.2           ucst  1321  24 ae1.0
8.8.8.8/32       user   0 3.3.3.2           ucst  1321  24 ae1.0
9.9.9.9/32       intf   0 9.9.9.9           locl  1280  1
10.10.10.10/32   user   0 3.3.3.2           ucst  1321  24 ae1.0
10.93.8.0/21     intf   0                               rslv  323  1 me0.0
10.93.8.0/32     dest   0 10.93.8.0         recv  321  1 me0.0
10.93.13.238/32  intf   0 10.93.13.238      locl  322  2
10.93.13.238/32  dest   0 10.93.13.238      locl  322  2
10.93.15.254/32  dest   0 0:12:f2:21:cf:0  ucst  333   5 me0.0
10.93.15.255/32  dest   0 10.93.15.255      bcst  320  1 me0.0
14.14.14.0/24    ifdn   0                               rslv  1319  1 ge-0/0/25.0
14.14.14.0/32    iddn   0 14.14.14.0        recv  1317  1 ge-0/0/25.0
14.14.14.2/32    user   0                               rjct   36   2
14.14.14.2/32    intf   0 14.14.14.2        locl  1318  2
14.14.14.2/32    iddn   0 14.14.14.2        locl  1318  2
14.14.14.255/32  iddn   0 14.14.14.255      bcst  1316  1 ge-0/0/25.0
224.0.0.0/4      perm   1                               mdsc   35   1
224.0.0.1/32     perm   0 224.0.0.1         mcst   31   3
224.0.0.5/32     user   1 224.0.0.5         mcst   31   3
255.255.255.255/32 perm   0                               bcst   32   1

```

### show route forwarding-table summary

```
user@switch> show route forwarding-table summary
```

```

Routing table: default.inet
Internet:
    user:          6 routes
    perm:          5 routes
    intf:          8 routes
    dest:         12 routes
    ifdn:          1 routes
    iddn:          3 routes

```

### show route forwarding-table extensive

```
user@switch> show route forwarding-table summary
```

```

Routing table: default.inet [Index 0]
Internet:

Destination: default
Route type: user
Route reference: 2
Route interface-index: 0

```

```

Flags: sent to PFE, rt nh decoupled
Nexthop: 0:12:f2:21:cf:0
Next-hop type: unicast          Index: 333      Reference: 5
Next-hop interface: me0.0

Destination: default
Route type: permanent
Route reference: 0              Route interface-index: 0
Flags: none
Next-hop type: reject          Index: 36       Reference: 2

Destination: 0.0.0.0/32
Route type: permanent
Route reference: 0              Route interface-index: 0
Flags: sent to PFE
Next-hop type: discard         Index: 34       Reference: 1

Destination: 2.2.2.0/24
Route type: interface
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Next-hop type: resolve         Index: 1309     Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.0/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Nexthop: 2.2.2.0
Next-hop type: receive         Index: 1307     Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.1/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Nexthop: 0:21:59:cc:89:c0
Next-hop type: unicast         Index: 1320     Reference: 1
Next-hop interface: ae0.0

Destination: 2.2.2.2/32
Route type: interface
Route reference: 0              Route interface-index: 0
Flags: sent to PFE
Nexthop: 2.2.2.2
Next-hop type: local           Index: 1308     Reference: 2

Destination: 2.2.2.2/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: none
Nexthop: 2.2.2.2
Next-hop type: local           Index: 1308     Reference: 2

Destination: 2.2.2.255/32
Route type: destination
Route reference: 0              Route interface-index: 66
Flags: sent to PFE
Nexthop: 2.2.2.255
Next-hop type: broadcast       Index: 1306     Reference: 1
Next-hop interface: ae0.0

```

**show route forwarding-table ccc**

```

user@switch> show route forwarding-table ccc ge-0/0/0.10
Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
ge-0/0/0.10      (CCC) user    0 3.3.3.2          Push 300112 1343  2 ae1.0

```

**show route forwarding-table family (MPLS)**

```

user@switch> show route forwarding-table family mpls

Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm    0
0                user    0          recv  49    3
1                user    0          recv  49    3
2                user    0          recv  49    3
299776           user    0          Pop   1334   2 ge-0/0/0.10
299792           user    0          Pop   1339   2 ge-0/0/0.14
299808           user    0          Pop   1341   2 ge-0/0/0.2
299824           user    0          Pop   1344   2 ge-0/0/0.11
299840           user    0          Pop   1345   2 ge-0/0/0.13
299856           user    0          Pop   1346   2 ge-0/0/0.18
299872           user    0          Pop   1347   2 ge-0/0/0.16
299888           user    0          Pop   1348   2 ge-0/0/0.7
299904           user    0          Pop   1349   2 ge-0/0/0.20
299920           user    0          Pop   1350   2 ge-0/0/0.19
299936           user    0          Pop   1351   2 ge-0/0/0.17
299952           user    0          Pop   1352   2 ge-0/0/0.9
299968           user    0          Pop   1353   2 ge-0/0/0.1
299984           user    0          Pop   1354   2 ge-0/0/0.12
300000           user    0          Pop   1355   2 ge-0/0/0.8
300016           user    0          Pop   1356   2 ge-0/0/0.4
300032           user    0          Pop   1357   2 ge-0/0/0.5
300048           user    0          Pop   1358   2 ge-0/0/0.3
300064           user    0          Pop   1359   2 ge-0/0/0.15
ge-0/0/0.1      (CCC) user    0 3.3.3.2          Push 300064 1340  2 ae1.0
ge-0/0/0.2      (CCC) user    0 3.3.3.2          Push 299872 1328  2 ae1.0
ge-0/0/0.3      (CCC) user    0 3.3.3.2          Push 299792 1323  2 ae1.0
ge-0/0/0.4      (CCC) user    0 3.3.3.2          Push 300016 1337  2 ae1.0
ge-0/0/0.5      (CCC) user    0 3.3.3.2          Push 299824 1325  2 ae1.0
ge-0/0/0.7      (CCC) user    0 3.3.3.2          Push 299920 1331  2 ae1.0
ge-0/0/0.8      (CCC) user    0 3.3.3.2          Push 299840 1326  2 ae1.0
ge-0/0/0.9      (CCC) user    0 3.3.3.2          Push 299888 1329  2 ae1.0
ge-0/0/0.10     (CCC) user    0 3.3.3.2          Push 300112 1343  2 ae1.0
ge-0/0/0.11     (CCC) user    0 3.3.3.2          Push 299776 1322  2 ae1.0
ge-0/0/0.12     (CCC) user    0 3.3.3.2          Push 299952 1333  2 ae1.0
ge-0/0/0.13     (CCC) user    0 3.3.3.2          Push 300096 1342  2 ae1.0
ge-0/0/0.14     (CCC) user    0 3.3.3.2          Push 299984 1335  2 ae1.0
ge-0/0/0.15     (CCC) user    0 3.3.3.2          Push 299936 1332  2 ae1.0
ge-0/0/0.16     (CCC) user    0 3.3.3.2          Push 299808 1324  2 ae1.0
ge-0/0/0.17     (CCC) user    0 3.3.3.2          Push 300000 1336  2 ae1.0
ge-0/0/0.18     (CCC) user    0 3.3.3.2          Push 300032 1338  2 ae1.0
ge-0/0/0.19     (CCC) user    0 3.3.3.2          Push 299904 1330  2 ae1.0
ge-0/0/0.20     (CCC) user    0 3.3.3.2          Push 299856 1327  2 ae1.0

```

**show route forwarding-table family (IPv6)**

```

user@switch> show route forwarding-table family inet6

```

```

Routing table: default.inet6
Internet6:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0                rjct  44    1
::/128           perm  0                dscd  42    1
ff00::/8         perm  0                mdsc  43    1
ff02::1/128      perm  0 ff02::1         mcst  39    1

```

```

Routing table: default-switch.inet6
Internet6:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0                rjct  530   1
::/128           perm  0                dscd  528   1
2:1::3a00/312    user  0                indr  131070 2
comp            572   1
2:1::3a82/320    user  0                indr  131071 3
comp            573   1
2:1::3af0/320    user  0                indr  131071 3
comp            573   1
2:1:0:ff00::/56  user  0                mdsc  529   2
ff00::/8         perm  0                mdsc  529   2
ff02::1/128      perm  0 ff02::1         mcst  526   1

```

```

Routing table: __master.anon__.inet6
Internet6:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0                rjct  554   1
::/128           perm  0                dscd  552   1
ff00::/8         perm  0                mdsc  553   1
ff02::1/128      perm  0 ff02::1         mcst  550   1

```

### show route forwarding-table label

```
user@switch> show route forwarding-table label 29976
```

```

Routing table: default.mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
299776           user  0                Pop   1334   2 ge-0/0/0.10

```

### show route forwarding-table matching

```
user@switch> show route forwarding-table matching 3
```

```

Routing table: default.inet
Internet:

```

### show route forwarding-table multicast

```
user@switch> show route forwarding-table multicast
```

```

Routing table: default.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
224.0.0.0/4       perm  1                mdsc  35    1
224.0.0.1/32      perm  0 224.0.0.1       mcst  31    3
224.0.0.5/32      user  1 224.0.0.5       mcst  31    3

```

```

Routing table: __master.anon__.inet
Internet:
Destination      Type RtRef Next hop      Type Index NhRef Netif
224.0.0.0/4       perm  0                mdsc  1289   1

```

```
224.0.0.1/32      perm      0 224.0.0.1      mcst  1285      1
```

```
Routing table: default.inet6
```

```
Internet6:
```

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
ff00::/8	perm	0		mdsc	43	1	
ff02::1/128	perm	0	ff02::1	mcst	39	1	

## show mpls interface

<b>List of Syntax</b>	<a href="#">Syntax on page 384</a> <a href="#">Syntax (EX Series Switches) on page 384</a>
<b>Syntax</b>	<pre>show mpls interface &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	show mpls interface
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces.
<b>Options</b>	<p><b>none</b>—Display information about MPLS-enabled interfaces.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display information about MPLS-enabled interfaces for the specified routing instance. If <b>instance-name</b> is omitted, information about MPLS-enabled interfaces is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Additional Information</b>	MPLS is enabled on an interface when the interface is configured with both the <b>set protocol mpls interface <i>interface-name</i></b> and <b>set interface <i>interface-name</i> unit 0 family mpls</b> statements.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls interface on page 385</a>
<b>Output Fields</b>	<p><a href="#">Table 38 on page 364</a> describes the output fields for the <b>show mpls interface</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 45: show mpls interface Output Fields**

Field Name	Field Description
<b>Interface</b>	Name of the interface.
<b>State</b>	State of the interface: <b>Up</b> or <b>Dn</b> (down).
<b>Administrative groups</b>	Administratively assigned colors of the link.



Table 45: show mpls interface Output Fields (*continued*)

Field Name	Field Description
Maximum labels	Maximum number of MPLS labels upon which MPLS can operate on a logical interface. This is configured using the <b>maximum-labels</b> statement at the [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] or the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family mpls] hierarchy levels.
Static protection revert time	Time (in seconds) that a static LSP must wait before traffic reverts from the bypass path to the original path. This is configured using the <b>protection-revert-time</b> statement at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels.
Always mark connection protection tlv	Enabled or Disabled: Enabled indicates that the <b>always-mark-connection-protection-tlv</b> statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. When this statement is configured, it marks all OAM traffic transiting this interface in preparation for switching the traffic to an alternate path based on the OAM functionality. To switch traffic to the bypass LSP, the <b>switch-away-lsps</b> statement must be configured.
Switch away lsps	Enabled or Disabled: Enabled indicates that the <b>switch-away-lsps</b> statement is configured at the [edit logical-systems <i>logical-system-name</i> protocols mpls interface <i>interface-name</i> static] or the [edit protocols mpls interface <i>interface-name</i> static] hierarchy levels. This enables you to switch an LSP away from a network node using a bypass LSP. This feature can be used in maintenance of active networks when a network device needs to be replaced without interrupting traffic passing through the network. The LSPs can be either static or dynamic.

## Sample Output

### show mpls interface

```
user@host> show mpls interface
```

```
Interface: ge-0/2/1.57
State: Up
Administrative group: <none>
Maximum labels: 5
Static protection revert time: 5 seconds
Always mark connection protection tlv: Disabled
Switch away lsps : Disabled
```

## show mpls lsp

**List of Syntax**   [Syntax on page 386](#)  
[Syntax \(EX Series Switches\) on page 386](#)

**Syntax**   `show mpls lsp`  
                   `<brief | detail | extensive | terse>`  
                   `<autobandwidth>`  
                   `<bidirectional | unidirectional>`  
                   `<bypass>`  
                   `<count-active-routes>`  
                   `<defaults>`  
                   `<descriptions>`  
                   `<down | up>`  
                   `<externally-controlled>`  
                   `<externally-provisioned>`  
                   `<logical-system (all | logical-system-name)>`  
                   `<lsp-type>`  
                   `<name name>`  
                   `<p2mp>`  
                   `<statistics>`  
                   `<transit>`

**Syntax (EX Series Switches)**   `show mpls lsp`  
                                   `<brief | detail | extensive | terse>`  
                                   `<bidirectional | unidirectional>`  
                                   `<bypass>`  
                                   `<descriptions>`  
                                   `<down | up>`  
                                   `<externally-controlled>`  
                                   `<externally-provisioned>`  
                                   `<lsp-type>`  
                                   `<name name>`  
                                   `<p2mp>`  
                                   `<statistics>`  
                                   `<transit>`

**Release Information**   Command introduced before Junos OS Release 7.4.  
                               **defaults** option added in Junos OS Release 8.5.  
                               Command introduced in Junos OS Release 9.5 for EX Series switches.  
                               Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description**   Display information about configured and active dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

**Options**   **none**—Display standard information about all configured and active dynamic MPLS LSPs.

**brief | detail | extensive | terse**—(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

**autobandwidth**—(Optional) Display automatic bandwidth information. This option is explained separately (see [show mpls lsp autobandwidth](#)).

**bidirectional | unidirectional**—(Optional) Display bidirectional or unidirectional LSP information, respectively.

**bypass**—(Optional) Display LSPs used for protecting other LSPs.

**count-active-routes**—(Optional) Display active routes for LSPs.

**defaults**—(Optional) Display the MPLS LSP default settings.

**descriptions**—(Optional) Display the MPLS label-switched path (LSP) descriptions. To view this information, you must configure the description statement at the **[edit protocol mpls lsp]** hierarchy level. Only LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

**down | up**—(Optional) Display only LSPs that are inactive or active, respectively.

**externally-controlled**—(Optional) Display the LSPs that are under the control of an external Path Computation Element (PCE).

**externally-provisioned**—(Optional) Display the LSPs that are generated dynamically and provisioned by an external Path Computation Element (PCE).

**instance *instance-name***—(Optional) Display MPLS LSP information for the specified instance. If *instance-name* is omitted, MPLS LSP information is displayed for the master instance.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

***lsp-type***—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

**name *name***—(Optional) Display information about the specified LSP or group of LSPs.

**p2mp**—(Optional) Display information about point-to-multipoint LSPs.

**statistics**—(Optional) (Ingress and transit routers only) Display accounting information about LSPs. Statistics are not available for LSPs on the egress routing device, because the penultimate routing device in the LSP sets the label to 0. Also, as the packet arrives at the egress routing device, the hardware removes its MPLS header and the packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.



**NOTE:** If a bypass LSP is configured for the primary static LSP, display cumulative statistics of packets traversing through the protected LSP and bypass LSP when traffic is re-optimized when the protected LSP link is restored. (Bypass LSPs are not supported on QFX Series switches.)

When used with the **bypass** option (**show mpls lsp bypass statistics**), display statistics for the traffic that flows only through the bypass LSP.

**transit**—(Optional) Display LSPs transiting this routing device.

**Required Privilege Level** view

**Related Documentation**

- [clear mpls lsp on page 317](#)
- [show mpls lsp autobandwidth on page 404](#)

**List of Sample Output**

- [show mpls lsp defaults on page 395](#)
- [show mpls lsp descriptions on page 395](#)
- [show mpls lsp detail on page 395](#)
- [show mpls lsp extensive on page 396](#)
- [show mpls lsp detail \(When Egress Protection Is in Effect During a Local Repair\) on page 397](#)
- [show mpls lsp extensive on page 397](#)
- [show mpls lsp ingress extensive on page 399](#)
- [show mpls lsp extensive \(automatic bandwidth adjustment enabled\) on page 400](#)
- [show mpls lsp bypass extensive on page 401](#)
- [show mpls lsp p2mp on page 401](#)
- [show mpls lsp p2mp detail on page 402](#)
- [show mpls lsp detail count-active-routes on page 402](#)
- [show mpls lsp statistics extensive on page 403](#)

**Output Fields** Table 46 on page 388 describes the output fields for the **show mpls lsp** command. Output fields are listed in the approximate order in which they appear.

**Table 46: show mpls lsp Output Fields**

Field Name	Field Description	Level of Output
<b>Ingress LSP</b>	Information about LSPs on the ingress routing device. Each session has one line of output.	All levels
<b>Egress LSP</b>	Information about the LSPs on the egress routing device. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. Each session has one line of output.	All levels
<b>Transit LSP</b>	Number of LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels

Table 46: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>P2MP name</b>	Name of the point-to-multipoint LSP. Dynamically generated P2MP LSPs used for VPLS flooding use dynamically generated P2MP LSP names. The name uses the format <i>identifier:vpls:router-id:routing-instance-name</i> . The <i>identifier</i> is automatically generated by Junos OS.	All levels
<b>P2MP branch count</b>	Number of destination LSPs the point-to-multipoint LSP is transmitting to.	All levels
<b>P</b>	An asterisk (*) under this heading indicates that the LSP is a primary path.	All levels
<b>address</b>	( <b>detail</b> and <b>extensive</b> ) Destination (egress routing device) of the LSP.	<b>detail extensive</b>
<b>To</b>	Destination (egress routing device) of the session.	<b>brief</b>
<b>From</b>	Source (ingress routing device) of the session.	<b>brief detail</b>
<b>State</b>	State of the LSP handled by this RSVP session: <b>Up</b> , <b>Dn</b> (down), or <b>Restart</b> .	<b>brief detail</b>
<b>Active Route</b>	Number of active routes (prefixes) installed in the forwarding table. For ingress LSPs, the forwarding table is the primary IPv4 table ( <b>inet.0</b> ). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table ( <b>mpls.0</b> ).	<b>detail extensive</b>
<b>Rt</b>	Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table ( <b>inet.0</b> ). For transit and egress RSVP sessions, the routing table is the primary MPLS table ( <b>mpls.0</b> ).	<b>brief</b>
<b>P</b>	Path. An asterisk (*) underneath this column indicates that the LSP is a primary path.	<b>brief</b>
<b>ActivePath</b>	(Ingress LSP) Name of the active path: <b>Primary</b> or <b>Secondary</b> .	<b>detail extensive</b>
<b>LSPname</b>	Name of the LSP.	<b>brief detail</b>
<b>Statistics</b>	Displays the number of packets and the number of bytes transmitted over the LSP. These counters are reset to zero whenever the LSP path is optimized (for example, during an automatic bandwidth allocation).	<b>extensive</b>
<b>Aggregate statistics</b>	Displays the number of packets and the number of bytes transmitted over the LSP. These counters continue to iterate even if the LSP path is optimized. You can reset these counters to zero using the <b>clear mpls lsp statistics</b> command.	<b>extensive</b>
<b>Packets</b>	Displays the number of packets transmitted over the LSP.	<b>brief extensive</b>
<b>Bytes</b>	Displays the number of bytes transmitted over the LSP.	<b>brief extensive</b>
<b>DiffServInfo</b>	Type of LSP: multiclass LSP ( <b>multiclass diffServ-TE LSP</b> ) or Differentiated-Services-aware traffic engineering LSP ( <b>diffServ-TE LSP</b> ).	<b>detail</b>

Table 46: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>LSPtype</b>	Type of LSP: <ul style="list-style-type: none"> <li>• <b>Static configured</b>—Static</li> <li>• <b>Dynamic configured</b>—Dynamic</li> <li>• <b>Externally controlled</b>—External path computing entity</li> </ul> Also indicates if the LSP is a <b>Penultimate hop popping</b> LSP or an <b>Ultimate hop popping</b> LSP.	<b>detail extensive</b>
<b>Bypass</b>	(Bypass LSP) Destination address (egress routing device) for the bypass LSP.	All levels
<b>LSPpath</b>	Indicates whether the RSVP session is for the primary or secondary LSP path. <b>LSPpath</b> can be either <b>primary</b> or <b>secondary</b> and can be displayed on the ingress, egress, and transit routing devices.	<b>detail</b>
<b>Bidir</b>	(GMPLS) The LSP allows data to travel in both directions between GMPLS devices.	All levels
<b>Bidirectional</b>	(GMPLS) The LSP allows data to travel both ways between GMPLS devices.	All levels
<b>FastReroute desired</b>	Fast reroute has been requested by the ingress routing device.	<b>detail</b>
<b>Link protection desired</b>	<b>detail</b>	
<b>Node/Link protection desired</b>	Link protection has been requested by the ingress routing device.	<b>detail</b>
<b>LoadBalance</b>	(Ingress LSP) CSPF load-balancing rule that was configured to select the LSP's path among equal-cost paths: <b>Most-fill</b> , <b>Least-fill</b> , or <b>Random</b> .	<b>detail extensive</b>
<b>Signal type</b>	Signal type for GMPLS LSPs. The signal type determines the peak data rate for the LSP: <b>DS0</b> , <b>DS3</b> , <b>STS-1</b> , <b>STM-1</b> , or <b>STM-4</b> .	All levels
<b>Encoding type</b>	LSP encoding type: <b>Packet</b> , <b>Ethernet</b> , <b>PDH</b> , <b>SDH/SONET</b> , <b>Lambda</b> , or <b>Fiber</b> .	All levels
<b>Switching type</b>	Type of switching on the links needed for the LSP: <b>Fiber</b> , <b>Lambda</b> , <b>Packet</b> , <b>TDM</b> , or <b>PSC-1</b> .	All levels
<b>GPID</b>	Generalized Payload Identifier (identifier of the payload carried by an LSP): <b>HDLC</b> , <b>Ethernet</b> , <b>IPv4</b> , <b>PPP</b> , or <b>Unknown</b> .	All levels
<b>Protection</b>	Configured protection capability desired for the LSP: <b>Extra</b> , <b>Enhanced</b> , <b>none</b> , <b>One plus one</b> , <b>One to one</b> , or <b>Shared</b> .	All levels
<b>Upstream label in</b>	(Bidirectional LSPs) Incoming label for reverse direction traffic for this LSP.	All levels
<b>Upstream label out</b>	(Bidirectional LSPs) Outgoing label for reverse direction traffic for this LSP.	All levels

Table 46: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Suggested label received</b>	(Bidirectional LSPs) Label the upstream node suggests to use in the Resv message that is sent.	All levels
<b>Suggested label sent</b>	(Bidirectional LSPs) Label the downstream node suggests to use in the Resv message that is returned.	All levels
<b>Autobandwidth</b>	(Ingress LSP) The LSP is performing autobandwidth allocation.	<b>detail extensive</b>
<b>MinBW</b>	(Ingress LSP) Configured minimum value of the LSP, in bps.	<b>detail extensive</b>
<b>MaxBW</b>	(Ingress LSP) Configured maximum value of the LSP, in bps.	<b>detail extensive</b>
<b>AdjustTimer</b>	(Ingress LSP) Configured value of the bandwidth adjustment timer, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	<b>detail extensive</b>
<b>Adjustment Threshold</b>	(Ingress LSP) Configured value for the <b>adjust-threshold</b> statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization.	<b>detail extensive</b>
<b>Time for Next Adjustment</b>	(Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.	<b>detail extensive</b>
<b>Time of Last Adjustment</b>	(Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.	<b>detail extensive</b>
<b>MaxAvgBW util</b>	(Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps.	<b>detail extensive</b>
<b>Overflow limit</b>	(Ingress LSP) Configured value of the threshold overflow limit.	<b>detail extensive</b>
<b>Overflow sample count</b>	(Ingress LSP) Current value for the overflow sample count.	<b>detail extensive</b>
<b>Bandwidth Adjustment in <i>nnn</i> second(s)</b>	(Ingress LSP) Current value of the bandwidth adjustment timer, indicating the amount of time remaining until the bandwidth adjustment will take place, in seconds.	<b>detail extensive</b>
<b>Underflow limit</b>	(Ingress LSP) Configured value of the threshold underflow limit.	<b>detail extensive</b>
<b>Underflow sample count</b>	(Ingress LSP) Current value for the underflow sample count.	<b>detail extensive</b>
<b>Underflow Max AvgBW</b>	(Ingress LSP) The highest sample bandwidth among the underflow samples recorded currently. This is the signaling bandwidth if an adjustment occurs because of an underflow.	<b>detail extensive</b>

Table 46: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Active path indicator</b>	(Ingress LSP) A value of * indicates that the path is active. The absence of * indicates that the path is not active. In the following example, "long" is the active path.  *Primary long Standby short	detail extensive
<b>Primary</b>	(Ingress LSP) Name of the primary path.	detail extensive
<b>Secondary</b>	(Ingress LSP) Name of the secondary path.	detail extensive
<b>Standby</b>	(Ingress LSP) Name of the path in standby mode.	detail extensive
<b>State</b>	(Ingress LSP) State of the path: <b>Up</b> or <b>Dn</b> (down).	detail extensive
<b>COS</b>	(Ingress LSP) Class-of-service value.	detail extensive
<b>Bandwidth per class</b>	(Ingress LSP) Active bandwidth for the LSP path for each MPLS class type, in bps.	detail extensive
<b>Priorities</b>	(Ingress LSP) Configured value of the setup priority and the hold priority respectively (the setup priority is displayed first), where 0 is the highest priority and 7 is the lowest priority. If you have not explicitly configured these values, the default values are displayed (7 for the setup priority and 0 for the hold priority).	detail extensive
<b>OptimizeTimer</b>	(Ingress LSP) Configured value of the optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	detail extensive
<b>SmartOptimizeTimer</b>	(Ingress LSP) Configured value of the smart optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	detail extensive
<b>Reoptimization in xxx seconds</b>	(Ingress LSP) Current value of the optimize timer, indicating the amount of time remaining until the path will be reoptimized, in seconds.	detail extensive
<b>Computed ERO (S [L] denotes strict [loose] hops)</b>	(Ingress LSP) Computed explicit route. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (S) or loose (L).	detail extensive
<b>CSPF metric</b>	(Ingress LSP) Constrained Shortest Path First metric for this path.	detail extensive



Table 46: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Received RRO</b>	<p>(Ingress LSP) Received record route. A series of hops, each with an address followed by a flag. (In most cases, the received record route is the same as the computed explicit route. If <b>Received RRO</b> is different from <b>Computed ERO</b>, there is a topology change in the network, and the route is taking a detour.) The following flags identify the protection capability and status of the downstream node:</p> <ul style="list-style-type: none"> <li>• <b>0x01</b>—Local protection available. The link downstream from this node is protected by a local repair mechanism. This flag can be set only if the Local protection flag was set in the <b>SESSION_ATTRIBUTE</b> object of the corresponding Path message.</li> <li>• <b>0x02</b>—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously).</li> <li>• <b>0x03</b>—Combination of <b>0x01</b> and <b>0x02</b>.</li> <li>• <b>0x04</b>—Bandwidth protection. The downstream routing device has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section.</li> <li>• <b>0x08</b>—Node protection. The downstream routing device has a backup path providing protection against link and node failure on the corresponding path section. If the downstream routing device can set up only a link-protection backup path, the <b>Local protection available</b> bit is set but the <b>Node protection</b> bit is cleared.</li> <li>• <b>0x09</b>—Detour is established. Combination of <b>0x01</b> and <b>0x08</b>.</li> <li>• <b>0x10</b>—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engine LSP. This flag indicates to the ingress legacy edge router (LER) of this LSP that it should be rerouted.</li> <li>• <b>0x20</b>—Node ID. Indicates that the address specified in the RRO's IPv4 or IPv6 sub-object is a node ID address, which refers to the router address or router ID. Nodes must use the same address consistently.</li> <li>• <b>0xb</b>—Detour is in use. Combination of <b>0x01</b>, <b>0x02</b>, and <b>0x08</b>.</li> </ul>	<b>detail extensive</b>
<b>Index number</b>	(Ingress LSP) Log entry number of each LSP path event. The numbers are in chronological descending order, with a maximum of 50 index numbers displayed.	<b>extensive</b>
<b>Date</b>	(Ingress LSP) Date of the LSP event.	<b>extensive</b>
<b>Time</b>	(Ingress LSP) Time of the LSP event.	<b>extensive</b>
<b>Event</b>	(Ingress LSP) Description of the LSP event.	<b>extensive</b>
<b>Created</b>	(Ingress LSP) Date and time the LSP was created.	<b>extensive</b>
<b>Resv style</b>	(Bypass) RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be <b>FF</b> (fixed filter), <b>SE</b> (shared explicit), or <b>WF</b> (wildcard filter).	<b>brief detail extensive</b>
<b>Labelin</b>	Incoming label for this LSP.	<b>brief detail</b>
<b>Labelout</b>	Outgoing label for this LSP.	<b>brief detail</b>

Table 46: show mpls lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>LSPname</b>	Name of the LSP.	<b>brief detail</b>
<b>Time left</b>	Number of seconds remaining in the lifetime of the reservation.	<b>detail</b>
<b>Since</b>	Date and time when the RSVP session was initiated.	<b>detail</b>
<b>Tspec</b>	Sender's traffic specification, which describes the sender's traffic parameters.	<b>detail</b>
<b>Port number</b>	Protocol ID and sender or receiver port used in this RSVP session.	<b>detail</b>
<b>PATH rcvfrom</b>	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this router, and number of packets received from the upstream neighbor.	<b>detail</b>
<b>PATH sentto</b>	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor, and number of packets sent to the downstream routing device.	<b>detail</b>
<b>RESV rcvfrom</b>	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor. The output in this field, which is consistent with that in the <b>PATH rcvfrom</b> field, indicates that the RSVP negotiation is complete.	<b>detail</b>
<b>Record route</b>	Recorded route for the session, taken from the record route object.	<b>detail</b>
<b>Soft preempt</b>	Number of soft preemptions that occurred on a path and when the last soft preemption occurred. Only successful soft preemptions are counted (those that actually resulted in a new path being used).	<b>detail</b>
<b>Soft preemption pending</b>	Path is in the process of being soft preempted. This display is removed once the ingress router has calculated a new path.	<b>detail</b>
<b>MPLS-TE LSP Defaults</b>	Default settings for MPLS traffic engineered LSPs: <ul style="list-style-type: none"> <li>• <b>LSP Holding Priority</b>—Determines the degree to which an LSP holds on to its session reservation after the LSP has been set up successfully.</li> <li>• <b>LSP Setup Priority</b>—Determines whether a new LSP that preempts an existing LSP can be established.</li> <li>• <b>Hop Limit</b>—Specifies the maximum number of routers the LSP can traverse (including the ingress and egress).</li> <li>• <b>Bandwidth</b>—Specifies the bandwidth in bits per second for the LSP.</li> <li>• <b>LSP Retry Timer</b>—Length of time in seconds that the ingress router waits between attempts to establish the primary path.</li> </ul>	<b>defaults</b>

The XML tag name of the **bandwidth** tag under the **auto-bandwidth** tag has been updated to **maximum-average-bandwidth**. You can see the new tag when you issue the **show mpls lsp extensive** command with the **| display xml** pipe option. If you have any scripts that use the **bandwidth** tag, ensure that they are updated to **maximum-average-bandwidth**.

## Sample Output

### show mpls lsp defaults

```
user@host> show mpls lsp defaults
MPLS-TE LSP Defaults
  LSP Holding Priority      0
  LSP Setup Priority       7
  Hop Limit                255
  Bandwidth                0
  LSP Retry Timer          30 seconds
```

### show mpls lsp descriptions

```
user@host> show mpls lsp descriptions
Ingress LSP: 3 sessions
To          LSP name          Description
10.0.0.195  to-sanjose                to-sanjose-desc
10.0.0.195  to-sanjose-other-desc      other-desc
Total 2 displayed, Up 2, Down 0
```

### show mpls lsp detail

```
user@host> show mpls lsp detail
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
  10.0.0.18 S 10.0.0.22 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
  20=Node-ID):
      10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 157, Since: Wed Jul 18 17:55:12 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 46128 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0
```

Transit LSP: 0 sessions  
Total 0 displayed, Up 0, Down 0

### show mpls lsp extensive

user@host> show mpls lsp extensive  
Ingress LSP: 1 sessions

```
192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  LSPtype: Static Configured, Ultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.0.0.18 10.0.0.22
11 Sep 20 15:54:35.032 Make-before-break: Switched to new instance
10 Sep 20 15:54:34.029 Record Route: 10.0.0.18 10.0.0.22
 9 Sep 20 15:54:34.029 Up
 8 Sep 20 15:54:20.271 Originate make-before-break call
 7 Sep 20 15:54:20.271 CSPF: computation result accepted 10.0.0.18 10.0.0.22

 6 Sep 20 15:52:10.247 Selected as active path
 5 Sep 20 15:52:10.246 Record Route: 10.0.0.18 10.0.0.22
 4 Sep 20 15:52:10.243 Up
 3 Sep 20 15:52:09.745 Originate Call
 2 Sep 20 15:52:09.745 CSPF: computation result accepted 10.0.0.18 10.0.0.22

 1 Sep 20 15:51:39.903 CSPF failed: no route toward 192.168.0.4
Created: Thu Sep 20 15:51:08 2012
Total 1 displayed, Up 1, Down 0
```

Egress LSP: 1 sessions

```
192.168.0.5
  From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 148, Since: Thu Sep 20 15:52:10 2012
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 49601 protocol 0
  PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 27 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0
```

Transit LSP: 0 sessions  
Total 0 displayed, Up 0, Down 0

**show mpls lsp detail (When Egress Protection Is in Effect During a Local Repair)**

```

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

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

Egress LSP: 1 sessions

192.168.0.5
  From: 192.168.0.4, LSPstate: Down, ActiveRoute: 0
  LSPname: E-D, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: -
  Resv style: 1 FF, Label in: 3, Label out: -
  Time left: 157, Since: Wed Jul 18 17:55:12 2012
  Tspecc: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 46128 protocol 0
Egress protection PLR as protector: In Use
PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 3 pkts
  Adspec: received MTU 1500
  PATH sentto: localclient
  RESV rcvfrom: localclient
  Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

**show mpls lsp extensive**

```

user@host> show mpls lsp extensive
Ingress LSP: 4 sessions

1.1.1.1
  From: 3.3.3.3, State: Up, ActiveRoute: 0, LSPname: m120b-to-mx960
  ActivePath: DEFAULT (primary)
  FastReroute desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary  DEFAULT                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 310)
10.0.35.5 S 10.0.15.1 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt

```

```

20=Node-ID):
    10.0.34.4(flag=1) 10.0.14.1
50 Sep 13 16:08:19.712 Record Route: 10.0.35.5(flag=1) 10.0.15.1
49 Sep 13 16:08:16.720 Record Route: 10.0.34.4(flag=1) 10.0.14.1
48 Sep 13 16:08:16.699 Fast-reroute Detour Up
47 Sep 13 16:08:13.702 Record Route: 10.0.34.4 10.0.14.1
46 Sep 13 16:08:13.702 Up
45 Sep 13 16:08:13.672 Originate make-before-break call
44 Sep 13 16:08:13.672 CSPF: computation result accepted 10.0.34.4 10.0.14.1

43 Sep 13 16:08:13.672 Selected as active path
42 Sep 13 16:08:13.672 Make-before-break: Switched to new instance
41 Sep 13 16:08:01.685 Pending path switchover, skip CSPF run[3 times]
40 Sep 13 16:06:33.910 Deselected as active
39 Sep 13 16:06:33.910 Pending path switchover, skip CSPF run

38 Sep 13 16:06:19.521 Record Route: 10.0.35.5 10.0.15.1
37 Sep 13 16:06:19.518 ResvTear received
36 Sep 13 16:06:19.518 Fast-reroute Detour Down
35 Sep 13 16:06:16.676 Record Route: 10.0.35.5(flag=1) 10.0.15.1
34 Sep 13 16:06:13.670 Record Route: 10.0.35.5 10.0.15.1
33 Sep 13 16:06:13.670 Up
32 Sep 13 16:06:13.569 Pending path switchover, skip CSPF run

31 Sep 13 16:06:13.569 CSPF: link down/deleted:
10.0.34.3(3.3.3.3:79)(m120-b-re1.00/3.3.3.3)->0.0.0.0(0.0.0.0:0)(m120-b-re1.04/0.0.0.0)

30 Sep 13 16:06:13.552 Pending path switchover, skip CSPF run

29 Sep 13 16:06:13.552 CSPF: link down/deleted:
0.0.0.0(0.0.0.0:0)(m120-b-re1.04/0.0.0.0)->0.0.0.0(4.4.4.4:0)(m10i-a-re0.00/4.4.4.4)

28 Sep 13 16:06:13.549 Originate make-before-break call
27 Sep 13 16:06:13.549 CSPF: computation result accepted 10.0.35.5 10.0.15.1

26 Sep 13 16:06:13.548 Tunnel local repaired
25 Sep 13 16:06:13.546 Record Route: 10.0.23.2 10.0.12.1
24 Sep 13 16:06:13.546 10.0.34.3: Tunnel local repaired
23 Sep 13 16:06:13.546 10.0.34.3: Down
22 Sep 13 16:03:46.842 Fast-reroute Detour Up
21 Sep 13 16:03:42.730 Record Route: 10.0.34.4(flag=1) 10.0.14.1
20 Sep 13 16:03:39.836 Selected as active path
19 Sep 13 16:03:39.834 Record Route: 10.0.34.4 10.0.14.1
18 Sep 13 16:03:39.834 Up
17 Sep 13 16:03:39.698 Originate Call
16 Sep 13 16:03:39.698 CSPF: computation result accepted 10.0.34.4 10.0.14.1

15 Sep 13 16:03:39.697 Clear Call
14 Sep 13 16:03:39.696 Deselected as active
13 Sep 13 16:03:37.837 Record Route: 10.0.34.4 10.0.14.1
12 Sep 13 16:03:32.829 Fast-reroute Detour Down
11 Sep 13 16:02:15.493 Record Route: 10.0.34.4(flag=1) 10.0.14.1
10 Sep 13 16:02:15.486 Fast-reroute Detour Up
9 Sep 13 16:02:12.468 Record Route: 10.0.34.4 10.0.14.1
8 Sep 13 16:02:07.460 Fast-reroute Detour Down
7 Sep 13 15:57:46.741 Fast-reroute Detour Up
6 Sep 13 15:57:40.768 Record Route: 10.0.34.4(flag=1) 10.0.14.1
5 Sep 13 15:57:37.761 Selected as active path
4 Sep 13 15:57:37.760 Record Route: 10.0.34.4 10.0.14.1
3 Sep 13 15:57:37.760 Up
2 Sep 13 15:57:37.733 Originate Call

```

```

1 Sep 13 15:57:37.733 CSPF: computation result accepted 10.0.34.4 10.0.14.1

Created: Fri Sep 13 15:57:38 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 4 sessions, 6 detours
Total 0 displayed, Up 0, Down 0

Transit LSP: 6 sessions, 1 detours

1.1.1.1
  From: 3.3.3.3, LSPstate: Up, ActiveRoute: 0
  LSPname: m120b-to-mx960, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 302288
  Resv style: 1 FF, Label in: 300416, Label out: 302288
  Time left: 147, Since: Fri Sep 13 16:08:16 2013
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 4 receiver 13955 protocol 0
  Detour branch from 10.0.34.4, to skip 1.1.1.1, Up
    Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
    Adspec: received MTU 1500
    Path MTU: received 0
    PATH rcvfrom: 10.0.34.4 (ge-4/3/7.0) 7 pkts
    Adspec: received MTU 1500 sent MTU 1500
    PATH sentto: 10.0.35.5 (ge-3/1/0.0) 7 pkts
    RESV rcvfrom: 10.0.35.5 (ge-3/1/0.0) 7 pkts
    Explicit route: 10.0.35.5 10.0.15.1
    Record route: 10.0.34.3 10.0.34.4 <self>10.0.35.5 10.0.15.1
  Label in: 300416, Label out: 302288
Total 1 displayed, Up 1, Down 0

```

### show mpls lsp ingress extensive

```

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

50.0.0.1
  From: 10.0.0.1, State: Up, ActiveRoute: 0, LSPname: test
  ActivePath: (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Priorities: 7 0
    OptimizeTimer: 300
    SmartOptimizeTimer: 180
    Reoptimization in 240 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 3)
    1.1.1.2 S 4.4.4.1 S 5.5.5.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
    20=Node-ID):
      1.1.1.2 4.4.4.1 5.5.5.2
    17 Aug 3 13:17:33.601 CSPF: computation result ignored, new path less avail
    bw[3 times]
    16 Aug 3 13:02:51.283 CSPF: computation result ignored, new path no benefit[2
    times]
    15 Aug 3 12:54:36.678 Selected as active path
    14 Aug 3 12:54:36.676 Record Route: 1.1.1.2 4.4.4.1 5.5.5.2
    13 Aug 3 12:54:36.676 Up
    12 Aug 3 12:54:33.924 Deselected as active

```

```

11 Aug 3 12:54:33.924 Originate Call
10 Aug 3 12:54:33.923 Clear Call
9 Aug 3 12:54:33.923 CSPF: computation result accepted 1.1.1.2 4.4.4.1
5.5.5.2
8 Aug 3 12:54:33.922 2.2.2.2: No Route toward dest
7 Aug 3 12:54:28.177 CSPF: computation result ignored, new path no benefit[4
times]
6 Aug 3 12:35:03.830 Selected as active path
5 Aug 3 12:35:03.828 Record Route: 2.2.2.2 3.3.3.2
4 Aug 3 12:35:03.827 Up
3 Aug 3 12:35:03.814 Originate Call
2 Aug 3 12:35:03.814 CSPF: computation result accepted 2.2.2.2 3.3.3.2
1 Aug 3 12:34:34.921 CSPF failed: no route toward 50.0.0.1
Created: Tue Aug 3 12:34:35 2010
Total 1 displayed, Up 1, Down 0

```

### show mpls lsp extensive (automatic bandwidth adjustment enabled)

```

user@host> show mpls lsp extensive
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPname: E-D
  ActivePath: (primary)
  Node/Link protection desired
  LSPtype: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Autobandwidth
  MinBW: 300bps, MaxBW: 1000bps, Dynamic MinBW: 1000bps
  Adjustment Timer: 300 secs AdjustThreshold: 25%
  Max AvgBW util: 963.739bps, Bandwidth Adjustment in 0 second(s).
  Min BW Adjust Interval: 1000, MinBW Adjust Threshold (in %): 50
  Overflow limit: 0, Overflow sample count: 0
  Underflow limit: 0, Underflow sample count: 9, Underflow Max AvgBW: 614.421bps

  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Priorities: 7 0
    Bandwidth: 1000bps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    192.168.0.6(flag=0x20) 10.0.0.18(Label=299792) 192.168.0.4(flag=0x20)
10.0.0.22(Label=3)
    12 Apr 30 10:25:17.024 Make-before-break: Switched to new instance
    11 Apr 30 10:25:16.023 Record Route: 192.168.0.6(flag=0x20)
10.0.0.18(Label=299792) 192.168.0.4(flag=0x20) 10.0.0.22(Label=3)
    10 Apr 30 10:25:16.023 Up
    9 Apr 30 10:25:16.023 Automatic Autobw adjustment succeeded: BW changes from
300 bps to 1000 bps
    8 Apr 30 10:25:15.946 Originate make-before-break call
    7 Apr 30 10:25:15.946 CSPF: computation result accepted 10.0.0.18 10.0.0.22

    6 Apr 30 10:16:42.891 Selected as active path
    5 Apr 30 10:16:42.891 Record Route: 192.168.0.6(flag=0x20)
10.0.0.18(Label=299776) 192.168.0.4(flag=0x20) 10.0.0.22(Label=3)
    4 Apr 30 10:16:42.890 Up
    3 Apr 30 10:16:42.828 Originate Call
    2 Apr 30 10:16:42.828 CSPF: computation result accepted 10.0.0.18 10.0.0.22

```



```

    1 Apr 30 10:16:14.064 CSPF: could not determine self[2 times]
Created: Tue Apr 30 10:15:16 2013
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

### show mpls lsp bypass extensive

```

user@host # show mpls lsp bypass extensive

Ingress LSP: 1 sessions

2.2.2.2
  From: 1.1.1.1, LSPstate: Up, ActiveRoute: 0
  LSPname: Bypass->1.1.2.2
  LSPtype: Static Configured
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 300032
  Resv style: 1 SE, Label in: -, Label out: 300032
  Time left: -, Since: Tue Dec 3 15:19:49 2013
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 1 receiver 55750 protocol 0
  Type: Bypass LSP
    Number of data route tunnel through: 1
    Number of RSVP session tunnel through: 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 1.1.5.2 (lt-1/2/0.15) 1221 pkts
  RESV rcvfrom: 1.1.5.2 (lt-1/2/0.15) 1221 pkts, Entropy label: No
  Explct route: 1.1.5.2 1.2.5.1
  Record route: <self> 1.1.5.2 1.2.5.1
+   4 Dec 3 15:19:49 Record Route: 1.1.5.2 1.2.5.1
+   3 Dec 3 15:19:49 Up
+   2 Dec 3 15:19:49 CSPF: computation result accepted
+   1 Dec 3 15:19:47 Originate Call
Total 1 displayed, Up 1, Down 0
Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
Transit LSP: 0 sessions

```

### show mpls lsp p2mp

```

user@host> show mpls lsp p2mp
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1
To          From          State Rt P ActivePath      LSPname
10.255.245.51 10.255.245.50 Up    0 * path1         p2mp-branch-1
P2MP name: p2mp-lsp2, P2MP branch count: 1
To          From          State Rt P ActivePath      LSPname
10.255.245.51 10.255.245.50 Up    0 * path1         p2mp-st-br1
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

```
Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

### show mpls lsp p2mp detail

```
user@host> show mpls lsp p2mp detail
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-branch-1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp1
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary path1 State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
  192.168.208.17 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    192.168.208.17
P2MP name: p2mp-lsp2, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-st-br1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp2
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary path1 State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
  192.168.208.17 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

    192.168.208.17
Total 2 displayed, Up 2, Down 0
```

### show mpls lsp detail count-active-routes

```
user@host> show mpls lsp detail count-active-routes
Ingress LSP: 1 sessions

213.119.192.2
  From: 156.154.162.128, State: Up, ActiveRoute: 1, LSPname: to-lahore
  ActivePath: (primary)
  LSPtype: Static Configured
  LoadBalance: Random
  Autobandwidth
  MinBW: 5Mbps MaxBW: 250Mbps
  AdjustTimer: 300 secs
  Max AvgBW util: 0bps, Bandwidth Adjustment in 102 second(s).
  Overflow limit: 0, Overflow sample count: 0
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary State: Up
    Priorities: 7 0
    Bandwidth: 5Mbps
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 4)
  10.252.0.177 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
    20=Node-ID):
```

```

10.252.0.177
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

### show mpls lsp statistics extensive

```

user@host> show mpls lsp statistics extensive
Ingress LSP: 1 sessions

192.168.0.4
  From: 192.168.0.5, State: Up, ActiveRoute: 0, LSPName: E-D
  Statistics: Packets 302, Bytes 28992
  Aggregate statistics: Packets 302, Bytes 28992
  ActivePath: (primary)
  LSPType: Static Configured, Penultimate hop popping
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  *Primary                               State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 30)
10.0.0.18 S 10.0.0.22 S
  Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
20=Node-ID):
    10.0.0.18 10.0.0.22
      6 Oct  3 11:18:28.281 Selected as active path
      5 Oct  3 11:18:28.281 Record Route:  10.0.0.18 10.0.0.22
      4 Oct  3 11:18:28.280 Up
      3 Oct  3 11:18:27.995 Originate Call
      2 Oct  3 11:18:27.995 CSPF: computation result accepted  10.0.0.18 10.0.0.22

      1 Oct  3 11:17:59.118 CSPF failed: no route toward 192.168.0.4[2 times]
  Created: Wed Oct  3 11:17:01 2012
Total 1 displayed, Up 1, Down 0

```

## show mpls lsp autobandwidth

<b>Syntax</b>	<code>show mpls lsp autobandwidth</code> <code>&lt;brief   detail   extensive&gt;</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 11.4. Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Display automatic bandwidth information for the LSP(s).
<b>Options</b>	<p><b>brief   detail   extensive</b> — (Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b> — (Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show mpls lsp on page 386</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show mpls lsp autobandwidth on page 405</a>
<b>Output Fields</b>	<a href="#">Table 47 on page 404</a> describes the output fields for the <b>show mpls lsp autobandwidth</b> command. Output fields are listed in the approximate order in which they appear.

**Table 47: show mpls lsp autobandwidth Output Fields**

Field Name	Field Description	Level of Output
<b>To</b>	Destination (egress routing device) of the session.	All Levels
<b>From</b>	Source (ingress routing device) of the session.	All Levels
<b>LSPname</b>	Name of the LSP.	All Levels
<b>Min BW</b>	(Ingress LSP) Configured minimum value of the LSP, in bps.	<b>detail extensive</b>
<b>Max BW</b>	(Ingress LSP) Configured maximum value of the LSP, in bps.	<b>detail extensive</b>
<b>Max AvgBW util</b>	(Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps.	<b>detail extensive</b>
<b>Overflow limit</b>	(Ingress LSP) Configured value of the threshold overflow limit.	<b>detail extensive</b>
<b>Overflow sample count</b>	(Ingress LSP) Current value for the overflow sample count.	<b>detail extensive</b>
<b>Underflow limit</b>	(Ingress LSP) Configured value of the threshold underflow limit.	<b>detail extensive</b>

Table 47: show mpls lsp autobandwidth Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Underflow sample count</b>	(Ingress LSP) Current value for the underflow sample count.	<b>detail extensive</b>
<b>Adjustment Timer</b>	(Ingress LSP) Configured value for the adjust-timer statement, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	<b>detail extensive</b>
<b>Adjustment Threshold</b>	(Ingress LSP) Configured value for the adjust-threshold statement. Specifies how sensitive the automatic bandwidth adjustment for an LSP is to changes in bandwidth utilization.	<b>detail extensive</b>
<b>Time for Next Adjustment</b>	(Ingress LSP) Time in seconds until the next automatic bandwidth adjustment sample is taken.	<b>detail extensive</b>
<b>Time of Last Adjustment</b>	(Ingress LSP) Date and time since the last automatic bandwidth adjustment was completed.	<b>detail extensive</b>
<b>Last BW</b>	Previous active bandwidth of the LSP.	<b>detail extensive</b>
<b>Last Requested BW</b>	Bandwidth requested in the previous automatic bandwidth adjustment.	<b>detail extensive</b>
<b>Last Signaled BW</b>	Bandwidth signaled in the previous automatic bandwidth adjustment.	<b>detail extensive</b>
<b>Highest Watermark BW</b>	Maximum bandwidth used by the LSP.	<b>detail extensive</b>
<b>Total AutoBw Adjustments</b>	Total number of attempts to adjust automatic bandwidth including failed and successful adjustments.	<b>detail extensive</b>
<b>Successful Adjustments</b>	Number of successful automatic bandwidth adjustments.	<b>detail extensive</b>
<b>Failed Adjustments</b>	Number of failed automatic bandwidth adjustments.	<b>detail extensive</b>

## Sample Output

### show mpls lsp autobandwidth

```

user@host> show mpls lsp autobandwidth extensive
To: 10.255.106.133,
From: 10.255.106.135, LSPname: r0-r1
Min BW: 100kbps, Max BW: 0bps, Max AvgBW util: 2.33249Mbps
Overflow limit: 0, Overflow sample count: 0
Underflow limit: 0, Underflow sample count: 0
Adjustment Timer: 300 sec, Adjustment Threshold: 0
Time for Next Adjustment: 23 sec, Time of Last Adjustment: Fri Jun 3 21:05:37
2011
Last BW: 100kbps, Last Requested BW: 2.2169Mbps, Last Signaled BW: 2.2169Mbps,
Highest Watermark BW: 2.33249Mbps
Total AutoBw Adjustments: 1, Successful Adjustments: 1, Failed Adjustments: 0

```



## show mpls path

<b>List of Syntax</b>	<a href="#">Syntax on page 407</a> <a href="#">Syntax (EX Series Switches) on page 407</a>
<b>Syntax</b>	<pre>show mpls path &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;path-name&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show mpls path &lt;path-name&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).
<b>Options</b>	<p><b>none</b>—Display standard information about all MPLS LSPs.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display the dynamic MPLS LSP for the specified instance. If <b><i>instance-name</i></b> is omitted, dynamic MPLS LSP for the master instance is displayed.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>path-name</i></b>—(Optional) Display information about the specified LSP only.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show mpls path on page 408</a>
<b>Output Fields</b>	<a href="#">Table 48 on page 407</a> describes the output fields for the <b>show mpls path</b> command. Output fields are listed in the approximate order in which they appear.

**Table 48: show mpls path Output Fields**

Field Name	Field Description
<b>Path name</b>	Information about ingress LSPs. Each path has one line of output.
<b>Address</b>	Addresses of the routing devices that form the LSP.
<b>Strict/loose address</b>	Whether the address is configured as a strict or loose address.

## Sample Output

### show mpls path

```
user@host> show mpls path
Path name      Address          Strict/loose address
p1             123.456.55.6    Strict
               123.456.1.6     Loose
p2             191.456.1.4     Strict
```



## show route table

<b>List of Syntax</b>	<a href="#">Syntax on page 409</a> <a href="#">Syntax (EX Series Switches and QFX Series Switches) on page 409</a>
<b>Syntax</b>	show route table <i>routing-table-name</i> <brief   detail   extensive   terse> <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switches and QFX Series Switches)</b>	show route table <i>routing-table-name</i> <brief   detail   extensive   terse>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches. Show route table evpn statement introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.
<b>Description</b>	Display the route entries in a particular routing table.
<b>Options</b>	<b>brief   detail   extensive   terse</b> —(Optional) Display the specified level of output.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.  <b><i>routing-table-name</i></b> —Display route entries for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the <b>show route table inet</b> command).
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show route summary</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show route table bgp.l2.vpn on page 420</a> <a href="#">show route table bgp.l3vpn.0 on page 420</a> <a href="#">show route table bgp.l3vpn.0 detail on page 420</a> <a href="#">show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 421</a> <a href="#">show route table bgp.evpn.0 on page 422</a> <a href="#">show route table evpna.evpn.0 on page 422</a> <a href="#">show route table inet.0 on page 422</a> <a href="#">show route table inet.3 on page 423</a> <a href="#">show route table inet6.0 on page 423</a> <a href="#">show route table inet6.3 on page 423</a> <a href="#">show route table inetflow detail on page 424</a> <a href="#">show route table l2circuit.0 on page 424</a> <a href="#">show route table mpls on page 424</a> <a href="#">show route table mpls extensive on page 425</a>

[show route table mpls.0 on page 425](#)  
[show route table mpls.0 detail \(PTX Series\) on page 425](#)  
[show route table mpls.0 extensive \(PTX Series\) on page 426](#)  
[show route table mpls.0 \(RSVP Route—Transit LSP\) on page 427](#)  
[show route table vpls\\_1 detail on page 427](#)  
[show route table vpn-a on page 427](#)  
[show route table vpn-a.mdt.0 on page 428](#)  
[show route table VPN-A detail on page 428](#)  
[show route table VPN-AB.inet.0 on page 429](#)  
[show route table VPN\\_blue.mvpn-inet6.0 on page 429](#)  
[show route table vrf1.mvpn.0 extensive on page 429](#)  
[show route table MVPN.mvpn.0 on page 430](#)  
[show route table inetflow detail on page 430](#)  
[show route table bgp.evpn.0 extensive |no-more \(EVPN\) on page 433](#)

**Output Fields** Table 49 on page 410 describes the output fields for the **show route table** command. Output fields are listed in the approximate order in which they appear.

**Table 49: show route table Output Fields**

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
Restart complete	<p>All protocols have restarted for this routing table.</p> <p>Restart state:</p> <ul style="list-style-type: none"> <li>• <b>Pending:</b><i>protocol-name</i>—List of protocols that have not yet completed graceful restart for this routing table.</li> <li>• <b>Complete</b>—All protocols have restarted for this routing table.</li> </ul> <p>For example, if the output shows-</p> <ul style="list-style-type: none"> <li>• LDP.inet.0 : 5 routes (4 active, 1 holddown, 0 hidden) Restart Pending: OSPF LDP VPN</li> </ul> <p>This indicates that <b>OSPF</b>, <b>LDP</b>, and <b>VPN</b> protocols did not restart for <b>LDP.inet.0</b> routing table.</p> <ul style="list-style-type: none"> <li>• vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden) Restart Complete</li> </ul> <p>This indicates that all protocols have restarted for <b>vpls_1.l2vpn.0</b> routing table.</p>
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	<p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> <li>• <b>active</b> (routes that are active)</li> <li>• <b>holddown</b> (routes that are in the pending state before being declared inactive)</li> <li>• <b>hidden</b> (routes that are not used because of a routing policy)</li> </ul>

Table 49: show route table Output Fields (*continued*)

Field Name	Field Description
<i>route-destination</i> (entry, announced)	<p>Route destination (for example:10.0.0.1/24). The <b>entry</b> value is the number of routes for this destination, and the <b>announced</b> value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> <li>• <b>MPLS-label</b> (for example, 80001).</li> <li>• <b>interface-name</b> (for example, ge-1/0/2).</li> <li>• <b>neighbor-address:control-word-status:encapsulation type:vc-id:source</b> (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> <li>• <b>neighbor-address</b>—Address of the neighbor.</li> <li>• <b>control-word-status</b>—Whether the use of the control word has been negotiated for this virtual circuit: <b>NoCtrlWord</b> or <b>CtrlWord</b>.</li> <li>• <b>encapsulation type</b>—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.</li> <li>• <b>vc-id</b>—Virtual circuit identifier.</li> <li>• <b>source</b>—Source of the advertisement: <b>Local</b> or <b>Remote</b>.</li> </ul> </li> <li>• <b>inclusive multicast Ethernet tag route</b>—Type of route destination represented by (for example, 3:100.100.100.10:100::0::10::100.100.100.10/384): <ul style="list-style-type: none"> <li>• <b>route distinguisher</b>—(8 octets) Route distinguisher (RD) must be the RD of the EVPN instance (EVI) that is advertising the NLRI.</li> <li>• <b>Ethernet tag ID</b>—(4 octets) Identifier of the Ethernet tag. Can set to 0 or to a valid Ethernet tag value.</li> <li>• <b>IP address length</b>—(1 octet) Length of IP address in bits.</li> <li>• <b>originating router's IP address</b>—(4 or 16 octets) Must set to the provider edge (PE) device's IP address. This address should be common for all EVIs on the PE device, and may be the PE device's loopback address.</li> </ul> </li> </ul>
label stacking	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> <li>• <b>S=0 route</b> indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed).</li> <li>• If there is no <b>S=</b> information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).</li> </ul>
[ <i>protocol, preference</i> ]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> <li>• <b>+</b>—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table.</li> <li>• <b>-</b>—A hyphen indicates the last active route.</li> <li>• <b>*</b>—An asterisk indicates that the route is both the active and the last active route. An asterisk before a <b>to</b> line indicates the best subpath to the route.</li> </ul> <p>In every routing metric except for the BGP <b>LocalPref</b> attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the <b>LocalPref</b> value in the <b>Preference2</b> field. For example, if the <b>LocalPref</b> value for Route 1 is 100, the <b>Preference2</b> value is -101. If the <b>LocalPref</b> value for Route 2 is 155, the <b>Preference2</b> value is -156. Route 2 is preferred because it has a higher <b>LocalPref</b> value and a lower <b>Preference2</b> value.</p>

Table 49: show route table Output Fields (*continued*)

Field Name	Field Description
Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see <a href="#">Table 50 on page 415</a> .
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word <b>Selected</b> . This field can also contain the following information: <ul style="list-style-type: none"> <li>Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.</li> <li>Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.</li> </ul>
Label-switched-path <i>lsp-path-name</i>	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be <b>pop</b> (where a label is removed from the top of the stack), <b>push</b> (where another label is added to the label stack), or <b>swap</b> (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See <a href="#">Table 51 on page 417</a> .

Table 49: show route table Output Fields (*continued*)

Field Name	Field Description
Local AS	AS number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.
Task	Name of the protocol that has added the route.
Announcement bits	<p>The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the KRT for installing the route into the Packet Forwarding Engine, to a resolve tree, a L2 VC, or even a VPN. For example, <i>n-Resolve inet</i> indicates that the specified route is used for route resolution for next hops found in the routing table.</p> <ul style="list-style-type: none"> <li><i>n</i>—An index used by Juniper Networks customer support only.</li> </ul>
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> <li><b>I</b>—IGP.</li> <li><b>E</b>—EGP.</li> <li><b>Recorded</b>—The AS path is recorded by the sample process (sampled).</li> <li><b>?</b>—Incomplete; typically, the AS path was aggregated.</li> </ul> <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> <li><b>[ ]</b>—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893.</li> <li><b>[ ]</b>—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path.</li> <li><b>{ }</b>—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order.</li> <li><b>( )</b>—Parentheses enclose a confederation.</li> <li><b>( [ ] )</b>—Parentheses and brackets enclose a confederation set.</li> </ul> <p><b>NOTE:</b> In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>

Table 49: show route table Output Fields (*continued*)

Field Name	Field Description
validation-state	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> <li>• <b>Invalid</b>—Indicates that the prefix is found, but either the corresponding AS received from the EBGp peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.</li> <li>• <b>Unknown</b>—Indicates that the prefix is not among the prefixes or prefix ranges in the database.</li> <li>• <b>Unverified</b>—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers.</li> <li>• <b>Valid</b>—Indicates that the prefix and autonomous system pair are found in the database.</li> </ul>
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See <a href="#">Table 52 on page 419</a> for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: <b>none</b> or <b>Site Down</b> .
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).

Table 49: show route table Output Fields (*continued*)

Field Name	Field Description
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag.
Accepted LongLivedStaleImport	<p>The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag.</p> <p>Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table</p>
ImportAccepted LongLivedStaleImport	<p>Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table</p> <p>The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.</p>
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 50 on page 415 describes all possible values for the Next-hop Types output field.

Table 50: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.

Table 50: Next-hop Types Output Field Values (*continued*)

Next-Hop Type	Description
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrtr)	Regular multicast next hop.
Router	<p>A specific node or set of nodes to which the routing device forwards packets that match the route prefix.</p> <p>To qualify as next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> <li>• Must not be a direct or local subnet for the routing device.</li> <li>• Must have a next hop that is directly connected to the routing device.</li> </ul>
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.



Table 51 on page 417 describes all possible values for the State output field. A route can be in more than one state (for example, <Active NoReadvrt Int Ext>).

**Table 51: State Output Field Values**

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGp path is available.

Table 51: State Output Field Values (*continued*)

Value	Description
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available.
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: <ul style="list-style-type: none"> <li>• The route is damped.</li> <li>• The route is rejected by an import policy.</li> <li>• The route is unresolved.</li> </ul>
Update source	Last tiebreaker is the lowest IP address value.

Table 52 on page 419 describes the possible values for the Communities output field.

**Table 52: Communities Output Field Values**

Value	Description
<i>area-number</i>	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
<b>bandwidth: local AS number:link-bandwidth-number</b>	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
<b>domain-id</b>	Unique configurable number that identifies the OSPF domain.
<b>domain-id-vendor</b>	Unique configurable number that further identifies the OSPF domain.
<i>link-bandwidth-number</i>	Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second).
<i>local AS number</i>	Local AS number: from 1 through 65,535.
<i>options</i>	1 byte. Currently this is only used if the route type is 5 or 7. Setting the least significant bit in the field indicates that the route carries a type 2 metric.
<b>origin</b>	(Used with VPNs) Identifies where the route came from.
<i>ospf-route-type</i>	1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be 0); 7 for NSSA routes; or 129 for sham link endpoint addresses.
<b>route-type-vendor</b>	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x8000. The format is <b>area-number:ospf-route-type:options</b> .
<b>rte-type</b>	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x0306. The format is <b>area-number:ospf-route-type:options</b> .
<b>target</b>	Defines which VPN the route participates in; <b>target</b> has the format <b>32-bit IP address:16-bit number</b> . For example, 10.19.0.0:100.
<b>unknown IANA</b>	Incoming IANA codes with a value between 0x1 and 0x7fff. This code of the BGP extended community attribute is accepted, but it is not recognized.
<b>unknown OSPF vendor community</b>	Incoming IANA codes with a value above 0x8000. This code of the BGP extended community attribute is accepted, but it is not recognized.

## Sample Output

### show route table bgp.l2vpn

```
user@host> show route table bgp.l2vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
    *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

### show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.15:100:10.255.71.17/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.15:200:10.255.71.18/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100021, Push 100011(top)
```

### show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.245.12:1
    Source: 10.255.245.12
    Next hop: 192.168.208.66 via fe-0/0/0.0, selected
    Label operation: Push 182449
    Protocol next hop: 10.255.245.12
    Push 182449
    Indirect next hop: 863a630 297
    State: <Active Int Ext>
    Local AS: 35 Peer AS: 35
    Age: 12:19 Metric2: 1
    Task: BGP_35.10.255.245.12+179
    Announcement bits (1): 0-BGP.0.0.0.0+179
    AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

    Communities: 2914:420 target:11111:1 origin:56:78
    VPN Label: 182449
    Localpref: 100
    Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.245.12:1
    Source: 10.255.245.12
    Next hop: 192.168.208.66 via fe-0/0/0.0, selected
```

```

Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 863a8f0 305
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100

```

### show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```
user@host> show route table bgp.rtarget.0
```

```

bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

100:100:100/96
    * [RTarget/5] 00:03:14
      Type Proxy
      for 10.255.165.103
      for 10.255.166.124
      Local

```

### show route table bgp.evpn.0

```

user@host> show route table bgp.evpn.0
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:100.100.100.2:100::0::00:26:88:5f:67:b0/304
    * [BGP/170] 11:00:05, localpref 100, from 100.100.100.2
      AS path: I, validation-state: unverified
      > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.2:100::0::00:51:51:51:51:51/304
    * [BGP/170] 11:00:05, localpref 100, from 100.100.100.2
      AS path: I, validation-state: unverified
      > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.3:100::0::00:52:52:52:52:52/304
    * [BGP/170] 10:59:58, localpref 100, from 100.100.100.3
      AS path: I, validation-state: unverified
      > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
2:100.100.100.3:100::0::a8:d0:e5:5b:01:c8/304
    * [BGP/170] 10:59:58, localpref 100, from 100.100.100.3
      AS path: I, validation-state: unverified
      > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
3:100.100.100.2:100::1000::100.100.100.2/304
    * [BGP/170] 11:00:16, localpref 100, from 100.100.100.2
      AS path: I, validation-state: unverified
      > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
3:100.100.100.2:100::2000::100.100.100.2/304
    * [BGP/170] 11:00:16, localpref 100, from 100.100.100.2
      AS path: I, validation-state: unverified
      > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1

```

### show route table evpna.evpn.0

```

user@host> show route table evpna.evpn.0
evpna.evpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

3:100.100.100.10:100::0::10::100.100.100.10/384
    * [EVPN/170] 01:37:09
      Indirect
3:100.100.100.2:100::2000::100.100.100.2/304
    * [EVPN/170] 01:37:12
      Indirect

```

### show route table inet.0

```

user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0
    * [Static/5] 00:51:57
      > to 111.222.5.254 via fxp0.0

```

```

1.0.0.1/32      *[Direct/0] 00:51:58
                 > via at-5/3/0.0
1.0.0.2/32      *[Local/0] 00:51:58
                 Local
12.12.12.21/32  *[Local/0] 00:51:57
                 Reject
13.13.13.13/32  *[Direct/0] 00:51:58
                 > via t3-5/2/1.0
13.13.13.14/32  *[Local/0] 00:51:58
                 Local
13.13.13.21/32  *[Local/0] 00:51:58
                 Local
13.13.13.22/32  *[Direct/0] 00:33:59
                 > via t3-5/2/0.0
127.0.0.1/32    [Direct/0] 00:51:58
                 > via lo0.0
111.222.5.0/24  *[Direct/0] 00:51:58
                 > via fxp0.0
111.222.5.81/32 *[Local/0] 00:51:58
                 Local

```

### show route table inet.3

```

user@host> show route table inet.3
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

22.0.0.5/32      *[LDP/9] 00:25:43, metric 10, tag 200
                  to 1.2.94.2 via lt-1/2/0.49
                  > to 1.2.3.2 via lt-1/2/0.23

```

### show route table inet6.0

```

user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0

```

### show route table inet6.3

```

user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
                  *[LDP/9] 00:00:22, metric 1
                  > via so-1/0/0.0
::10.255.245.196/128
                  *[LDP/9] 00:00:08, metric 1
                  > via so-1/0/0.0, Push 100008

```

**show route table inetflow detail**

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP    Preference: 170/-101
            Next-hop reference count: 2
            State: <Active Ext>
            Local AS: 65002 Peer AS: 65000
            Age: 4
            Task: BGP_65000.10.12.99.5+3792
            Announcement bits (1): 0-Flow
            AS path: 65000 I
            Communities: traffic-rate:0:0
            Validation state: Accept, Originator: 10.12.99.5
            Via: 10.12.44.0/24, Active
            Localpref: 100
            Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
    *Flow    Preference: 5
            Next-hop reference count: 2
            State: <Active>
            Local AS: 65002
            Age: 6:30
            Task: RT Flow
            Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
            AS path: I
            Communities: 1:1

```

**show route table l2circuit.0**

```

user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
    *[L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
    *[LDP/9] 00:50:14
    Discard
10.1.1.195:CtrlWord:1:2:Local/96
    *[L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
    *[LDP/9] 00:50:14
    Discard

```

**show route table mpls**

```

user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 00:13:55, metric 1
           Receive
1          *[MPLS/0] 00:13:55, metric 1
           Receive

```



```

2          *[MPLS/0] 00:13:55, metric 1
           Receive
1024       *[VPN/0] 00:04:18
           to table red.inet.0, Pop

```

### show route table mpls extensive

```

user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
    *LDP   Preference: 9
           Next hop: via so-1/0/0.0, selected
           Pop
           State: <Active Int>
           Age: 29:50      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
           Prefixes bound to route: 10.0.0.194/32

```

### show route table mpls.0

```

user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0          *[MPLS/0] 00:45:09, metric 1
           Receive
1          *[MPLS/0] 00:45:09, metric 1
           Receive
2          *[MPLS/0] 00:45:09, metric 1
           Receive
100000     *[L2VPN/7] 00:43:04
           > via so-0/1/0.1, Pop
100001     *[L2VPN/7] 00:43:03
           > via so-0/1/0.2, Pop      Offset: 4
100002     *[LDP/9] 00:43:22, metric 1
           via so-0/1/2.0, Pop
           > via so-0/1/3.0, Pop
100002(S=0) *[LDP/9] 00:43:22, metric 1
           via so-0/1/2.0, Pop
           > via so-0/1/3.0, Pop
100003     *[LDP/9] 00:43:22, metric 1
           > via so-0/1/2.0, Swap 100002
           via so-0/1/3.0, Swap 100002
100004     *[LDP/9] 00:43:16, metric 1
           via so-0/1/2.0, Swap 100049
           > via so-0/1/3.0, Swap 100049
so-0/1/0.1 *[L2VPN/7] 00:43:04
           > via so-0/1/2.0, Push 100001, Push 100049(top)
           via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2 *[L2VPN/7] 00:43:03
           via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
           > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4

```

### show route table mpls.0 detail (PTX Series)

```

user@host> show route table mpls.0 detail
ge-0/0/2.600 (1 entry, 1 announced)
    *L2VPN Preference: 7
           Next hop type: Indirect

```

```

Address: 0x9438f34
Next-hop reference count: 2
Next hop type: Router, Next hop index: 567
Next hop: 3.0.0.1 via ge-0/0/1.0, selected
Label operation: Push 299808
Label TTL action: prop-ttl
Load balance label: Label 299808:None;
Session Id: 0x1
Protocol next hop: 10.255.255.1
Label operation: Push 299872 Offset: 252
Label TTL action: no-prop-ttl
Load balance label: Label 299872:Flow label PUSH;
Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
State: <Active Int>
Age: 21          Metric2: 1
Validation State: unverified
Task: Common L2 VC
Announcement bits (2): 0-KRT 2-Common L2 VC
AS path: I

```

#### show route table mpls.0 extensive (PTX Series)

```

user@host> show route table mpls.0 extensive
ge-0/0/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel ge-0/0/2.600.0      /32 -> {composite(570)}
    *L2VPN Preference: 7
      Next hop type: Indirect
      Address: 0x9438f34
      Next-hop reference count: 2
      Next hop type: Router, Next hop index: 567
      Next hop: 3.0.0.1 via ge-0/0/1.0, selected
      Label operation: Push 299808
      Label TTL action: prop-ttl
      Load balance label: Label 299808:None;
      Session Id: 0x1
      Protocol next hop: 10.255.255.1
      Label operation: Push 299872 Offset: 252
      Label TTL action: no-prop-ttl
      Load balance label: Label 299872:Flow label PUSH;
      Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
      Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
      State: <Active Int>
      Age: 47          Metric2: 1
      Validation State: unverified
      Task: Common L2 VC
      Announcement bits (2): 0-KRT 2-Common L2 VC
      AS path: I
      Composite next hops: 1
        Protocol next hop: 10.255.255.1 Metric: 1
        Label operation: Push 299872 Offset: 252
        Label TTL action: no-prop-ttl
        Load balance label: Label 299872:Flow label PUSH;
        Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
        Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
        Indirect path forwarding next hops: 1
          Next hop type: Router
          Next hop: 3.0.0.1 via ge-0/0/1.0
          Session Id: 0x1
          10.255.255.1/32 Originating RIB: inet.3

```

```

Metric: 1
Forwarding nexthops: 1
Node path count: 1
Nexthop: 3.0.0.1 via ge-0/0/1.0

```

### show route table mpls.0 (RSVP Route—Transit LSP)

In the sample output, the 1 in [RSVP/7/1] indicates the secondary preference value. The secondary preference value becomes significant when multiple RSVP LSPs of different types are signaled to the destination. The possible values of RSVP secondary preferences are:

1—Normal Point-to-Point RSVP-TE LSP

2—Point-to-Multipoint (P2MP) RSVP-TE LSP

3—Dynamic RSVP-TE LSP

```
user@host> show route table mpls.0
```

```

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

0          *[MPLS/0] 00:37:31, metric 1
            Receive
1          *[MPLS/0] 00:37:31, metric 1
            Receive
2          *[MPLS/0] 00:37:31, metric 1
            Receive
13         *[MPLS/0] 00:37:31, metric 1
            Receive
300352     *[RSVP/7/1] 00:08:00, metric 1
            > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0) *[RSVP/7/1] 00:08:00, metric 1
            > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384     *[RSVP/7/2] 00:05:20, metric 1
            > to 8.64.1.106 via ge-1/0/0.0, Pop
300384(S=0) *[RSVP/7/2] 00:05:20, metric 1
            > to 8.64.1.106 via ge-1/0/0.0, Pop

```

### show route table vpls\_1 detail

```
user@host> show route table vpls_1 detail
```

```

vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

```

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF

```

### show route table vpn-a

```
user@host> show route table vpn-a
```

```

vpn-a.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
    *[VPN/7] 05:48:27
    Discard
192.168.24.1:1:2:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

#### show route table vpn-a.mdt.0

```

user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
    *[MVPN/70] 01:23:05, metric2 1
    Indirect
1:1:1:10.255.14.218:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
    AS path: I
    > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
    AS path: I
    > via so-0/0/1.0, label-switched-path r0-to-r2

```

#### show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
    *BGP Preference: 170/-101
    Route Distinguisher: 10.255.179.13:200
    Next hop type: Indirect
    Next-hop reference count: 5
    Source: 10.255.179.13
    Next hop type: Router, Next hop index: 732
    Next hop: 10.39.1.14 via fe-0/3/0.0, selected
    Label operation: Push 299824, Push 299824(top)
    Protocol next hop: 10.255.179.13
    Push 299824
    Indirect next hop: 8f275a0 1048574
    State: (Secondary Active Int Ext)
    Local AS: 1 Peer AS: 1
    Age: 3:41:06 Metric: 1 Metric2: 1
    Task: BGP_1.10.255.179.13+64309
    Announcement bits (2): 0-KRT 1-BGP RT Background
    AS path: I
    Communities: target:1:200 rte-type:0.0.0.0:1:0
    Import Accepted
    VPN Label: 299824 TTL Action: vrf-ttl-propagate
    Localpref: 100
    Router ID: 10.255.179.13
    Primary Routing Table bgp.13vpn.0

```

**show route table VPN-AB.inet.0**

```

user@host> show route table VPN-AB.inet.0
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.0/30      *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0
10.39.1.4/30      *[Direct/0] 00:08:42
                  > via so-5/1/0.0
10.39.1.6/32      *[Local/0] 00:08:46
                  Local
10.255.71.16/32   *[Static/5] 00:07:24
                  > via so-2/0/0.0
10.255.71.17/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.18/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32 *[BGP/170] 00:08:35, localpref 100
                  AS path: 2 I
                  > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32 *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0

```

**show route table VPN\_blue.mvpn-inet6.0**

```

user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65535:10.255.2.202/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                  AS path: I
                  > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
                  AS path: I
                  > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
                  *[MVPN/70] 00:57:23, metric2 1
                  Indirect
5:10.255.2.202:65535:128::192.168.90.2:128:ffff::1/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                  AS path: I
                  > via so-0/1/3.0
6:10.255.2.203:65535:65000:128::10.12.53.12:128:ffff::1/432
                  *[PIM/105] 00:02:37
                  Multicast (IPv6)
7:10.255.2.202:65535:65000:128::192.168.90.2:128:ffff::1/432
                  *[MVPN/70] 00:02:37, metric2 1
                  Indirect

```

**show route table vrf1.mvpn.0 extensive**

```

user@host> show route table vrf1.mvpn.0 extensive
1:10.255.50.77:1:10.255.50.77/240 (1 entry, 1 announced)
    *MVPN    Preference: 70

```

```

PMSI: Flags 0x0: Label 0: RSVP-TE:
Session_13[10.255.50.77:0:25624:10.255.50.77]
  Next hop type: Indirect
  Address: 0xbb2c944
  Next-hop reference count: 360
  Protocol next hop: 10.255.50.77
  Indirect next hop: 0x0 - INH Session ID: 0x0
  State: <Active Int Ext>
  Age: 53:03      Metric2: 1
  Validation State: unverified
  Task: mvpn global task
  Announcement bits (3): 0-PIM.vrf1 1-mvpn global task 2-rt-export

AS path: I

```

### show route table MVPN.mvpn.0

Starting in Junos OS Release 15.1, multicast routes on the locally originated type 7 customer multicast routes are added exclusively by PIM. The functionality of the BGP-MVPN service (which, internally, depends on contributions of state from both the MVPN and PIM protocol components of Junos OS) remains unchanged. MVPN, however, no longer appears as the originator of the locally advertised route. Routes advertised by remote PEs are, as usual, always learned locally from their respective [BGP/...] protocol.

```

user@host> show route table MVPN.mvpn.0
MVPN.mvpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

7:10.255.2.202:65535:65000:128:::192.168.90.2:128:ffff::1/432
    *[PIM/70] 00:02:37, metric2 1
    Indirect
5:100:32:192.168.1.9:32:239.1.1.1/240
    *[PIM/105] 01:51:21
    Multicast (IPv4)
7:100:1:100.32.192.168.5:32:237.1.1.1/240
    *[PIM/105] 01:51:21
    Multicast (IPv4)

```

### show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP      Preference: 170/-101
    Next-hop reference count: 2
    State: <Active Ext>
    Local AS: 65002 Peer AS: 65000
    Age: 4
    Task: BGP_65000.10.12.99.5+3792
    Announcement bits (1): 0-Flow
    AS path: 65000 I
    Communities: traffic-rate:0:0
    Validation state: Accept, Originator: 10.12.99.5
    Via: 10.12.44.0/24, Active
    Localpref: 100
    Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
    *Flow      Preference: 5
    Next-hop reference count: 2

```

```

State: <Active>
Local AS: 65002
Age: 6:30
Task: RT Flow
Announcement bits (2): 0-Flow 1-BGP.0.0.0+179
AS path: I
Communities: 1:1

user@PE1> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1.1.1.4/96 AD
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/1.5
1.1.1.2:100:1:0/96 MH
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1:0/96 MH
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
    *[VPLS/7] 1d 03:11:02, metric2 1
    > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
    *[LDP/9] 1d 03:11:02
    Discard

user@host> show route table red extensive
red.inet.0: 364481 destinations, 714087 routes (364480 active, 48448 holddown, 1
hidden)
22.0.0.0/32 (3 entries, 1 announced)
    State: <OnList CalcForwarding>
TSI:
KRT in-kernel 22.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

    Nexthop: Self
    AS path: [2] I
    Communities: target:2:1
Path 22.0.0.0 from 2.3.0.0 Vector len 4. Val: 1
    @BGP Preference: 170/-1
    Route Distinguisher: 2:1
    Next hop type: Indirect
    Address: 0x258059e4
    Next-hop reference count: 2
    Source: 2.2.0.0
    Next hop type: Router
    Next hop: 10.1.1.1 via ge-1/1/9.0, selected
    Label operation: Push 707633
    Label TTL action: prop-ttl
    Session Id: 0x17d8
    Protocol next hop: 2.2.0.0
    Push 16
    Composite next hop: 0x25805988 - INH Session ID: 0x193c
    Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
    State: <Secondary Active Int Ext ProtectionPath ProtectionCand>

```

```

Local AS:      2 Peer AS:      2
Age: 23        Metric2: 35
Validation State: unverified
Task: BGP_2.2.2.0.0+34549
AS path: I
Communities: target:2:1
Import Accepted
VPN Label: 16
Localpref: 0
Router ID: 2.2.0.0
Primary Routing Table bgp.13vpn.0
Composite next hops: 1
    Protocol next hop: 2.2.0.0 Metric: 35
    Push 16
    Composite next hop: 0x25805988 - INH Session ID: 0x193c
    Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
    Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.1.1.1 via ge-1/1/9.0
        Session Id: 0x17d8
    2.2.0.0/32 Originating RIB: inet.3
    Metric: 35                      Node path count: 1
    Forwarding nexthops: 1
        Nexthop: 10.1.1.1 via ge-1/1/9.0
BGP Preference: 170/-1
Route Distinguisher: 2:1
Next hop type: Indirect
Address: 0x9347028
Next-hop reference count: 3
Source: 2.3.0.0
Next hop type: Router, Next hop index: 702
Next hop: 10.1.4.2 via ge-1/0/0.0, selected
Label operation: Push 634278
Label TTL action: prop-ttl
Session Id: 0x17d9
Protocol next hop: 2.3.0.0
Push 16
Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da
State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

Inactive reason: Not Best in its group - IGP metric
Local AS:      2 Peer AS:      2
Age: 3:34      Metric2: 70
Validation State: unverified
Task: BGP_2.2.3.0.0+32805
Announcement bits (2): 0-KRT 1-BGP_RT_Background
AS path: I
Communities: target:2:1
Import Accepted
VPN Label: 16
Localpref: 0
Router ID: 2.3.0.0
Primary Routing Table bgp.13vpn.0
Composite next hops: 1
    Protocol next hop: 2.3.0.0 Metric: 70
    Push 16
    Composite next hop: 0x93463a0 1048575 INH Session ID:
0x17da
    Indirect next hop: 0x91e8800 1048574 INH Session ID:
0x17da

```



```

        Indirect path forwarding next hops: 1
            Next hop type: Router
            Next hop: 10.1.4.2 via ge-1/0/0.0
            Session Id: 0x17d9
        2.3.0.0/32 Originating RIB: inet.3
            Metric: 70
            Node path count: 1
            Forwarding nexthops: 1
            Nexthop: 10.1.4.2 via ge-1/0/0.0
#Multipath Preference: 255
    Next hop type: Indirect
    Address: 0x24afca30
    Next-hop reference count: 1
    Next hop type: Router
    Next hop: 10.1.1.1 via ge-1/1/9.0, selected
    Label operation: Push 707633
    Label TTL action: prop-ttl
    Session Id: 0x17d8
    Next hop type: Router, Next hop index: 702
    Next hop: 10.1.4.2 via ge-1/0/0.0
    Label operation: Push 634278
    Label TTL action: prop-ttl
    Session Id: 0x17d9
    Protocol next hop: 2.2.0.0
    Push 16
    Composite next hop: 0x25805988 - INH Session ID: 0x193c
    Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

    Protocol next hop: 2.3.0.0
    Push 16
    Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
    Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da Weight

0x4000
    State: <ForwardingOnly Int Ext>
    Inactive reason: Forwarding use only
    Age: 23      Metric2: 35
    Validation State: unverified
    Task: RT
    AS path: I
    Communities: target:2:1

```

### show route table bgp.evpn.0 extensive |no-more (EVPN)

```

show route table bgp.evpn.0 extensive | no-more
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
2:1000:10::100::00:aa:aa:aa:aa:aa/304 (1 entry, 0 announced)
    *BGP      Preference: 170/-101
              Route Distinguisher: 1000:10
              Next hop type: Indirect
              Address: 0x9420fd0
              Next-hop reference count: 12
              Source: 1.2.3.4
              Protocol next hop: 1.2.3.4
              Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    State: Local AS: 17 Peer AS:17 Age:21:12 Metric2:1 Validation State:
unverified
              Task: BGP_17.1.2.3.4+50756
              AS path: I
              Communities: target:1111:8388708 encapsulation0:0:0:0:3
              Import Accepted
              Route Label: 100
              ESI: 00:00:00:00:00:00:00:00:00:00

```

```

Localpref: 100
Router ID: 1.2.3.4
Secondary Tables: default-switch.evpn.0
Indirect next hops: 1
    Protocol next hop: 1.2.3.4 Metric: 1
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.10.10.1 via xe-0/0/1.0
        Session Id: 0x2
    1.2.3.4/32 Originating RIB: inet.0
        Metric: 1                      Node path count: 1
        Forwarding nexthops: 2
        Nexthop: 10.92.78.102 via em0.0

2:1000:10::200::00:bb:bb:bb:bb:bb/304 (1 entry, 0 announced)
    *BGP    Preference: 170/-101
            Route Distinguisher: 1000:10
            Next hop type: Indirect
            Address: 0x9420fd0
            Next-hop reference count: 12
            Source: 1.2.3.4
            Protocol next hop: 1.2.3.4
            Indirect next hop: 0x2 no-forward INH Session ID: 0x0
            State: Local AS:17 Peer AS:17 Age:19:43 Metric2:1 Validation
State:unverified
            Task: BGP_17.1.2.3.4+50756
            AS path: I
            Communities: target:2222:22 encapsulation0:0:0:0:3
            Import Accepted
            Route Label: 200
            ESI: 00:00:00:00:00:00:00:00:00:00
            Localpref: 100
            Router ID: 1.2.3.4
            Secondary Tables: default-switch.evpn.0
            Indirect next hops: 1
                Protocol next hop: 1.2.3.4 Metric: 1
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                Indirect path forwarding next hops: 1
                    Next hop type: Router
                    Next hop: 10.10.10.1 via xe-0/0/1.0
                    Session Id: 0x2
                1.2.3.4/32 Originating RIB: inet.0
                    Metric: 1                      Node path count: 1
                    Forwarding nexthops: 2
                    Nexthop: 10.92.78.102 via em0.0

2:1000:10::300::00:cc:cc:cc:cc:cc/304 (1 entry, 0 announced)
    *BGP    Preference: 170/-101
            Route Distinguisher: 1000:10
            Next hop type: Indirect
            Address: 0x9420fd0
            Next-hop reference count: 12
            Source: 1.2.3.4
            Protocol next hop: 1.2.3.4
            Indirect next hop: 0x2 no-forward INH Session ID: 0x0
            State: Local AS:17 Peer AS:17 Age:17:21 Metric2:1 Validation State:
unverified Task: BGP 17,1,2,3,4+50756
            AS path: I
            Communities: target:3333:33 encapsulation0:0:0:0:3
            Import Accepted

```

```

Route Label: 300
ESI: 00:00:00:00:00:00:00:00:00
Localpref: 100
Router ID: 1.2.3.4
Secondary Tables: default-switch.evpn.0
Indirect next hops: 1
    Protocol next hop: 1.2.3.4 Metric: 1
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.10.10.1 via xe-0/0/1.0
        Session Id: 0x2
    1.2.3.4/32 Originating RIB: inet.0
        Metric: 1                      Node path count: 1
        Forwarding nexthops: 2
        Nexthop: 10.92.78.102 via em0.0

3:1000:10::100::1.2.3.4/304 (1 entry, 0 announced)
*BGP   Preference: 170/-101
      Route Distinguisher: 1000:10
      PMSI: Flags 0x0: Label 100: Type INGRESS-REPLICATION 1.2.3.4
      Next hop type: Indirect
      Address: 0x9420fd0
      Next-hop reference count: 12
      Source: 1.2.3.4
      Protocol next hop: 1.2.3.4
      Indirect next hop: 0x2 no-forward INH Session ID: 0x0
      State: Local AS:17 Peer AS:17 Age:37:01 Metric2:1 Validation State:
unverified Task: BGP 17.1.2.3.4+50756
      AS path: I
      Communities: target:1111:8388708 encapsulation0:0:0:0:3
      Import Accepted
      Localpref: 100
      Router ID: 1.2.3.4
      Secondary Tables: default-switch.evpn.0
      Indirect next hops: 1
          Protocol next hop: 1.2.3.4 Metric: 1
          Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          Indirect path forwarding next hops: 1
              Next hop type: Router
              Next hop: 10.10.10.1 via xe-0/0/1.0
              Session Id: 0x2
          1.2.3.4/32 Originating RIB: inet.0
              Metric: 1                      Node path count: 1
              Forwarding nexthops: 2
              Nexthop: 10.92.78.102 via em0.0

3:1000:10::200::1.2.3.4/304 (1 entry, 0 announced)
*BGP   Preference: 170/-101
      Route Distinguisher: 1000:10
      PMSI: Flags 0x0: Label 200: Type INGRESS-REPLICATION 1.2.3.4
      Next hop type: Indirect
      Address: 0x9420fd0
      Next-hop reference count: 12
      Source: 1.2.3.4
      Protocol next hop: 1.2.3.4
      Indirect next hop: 0x2 no-forward INH Session ID: 0x0
      State: Local AS: 17 Peer AS: 17 Age:35:22 Metric2:1 Validation
State:unverified Task: BGP 17.1.2.3.4+50756
      AS path:I Communities: target:2222:22 encapsulation):0:0:0:0:3

```

```

Import Accepted
  Localpref: 100
  Router ID: 1.2.3.4
  Secondary Tables: default-switch.evpn.0
  Indirect next hops: 1
    Protocol next hop: 1.2.3.4 Metric: 1
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    Indirect path forwarding next hops: 1
      Next hop type: Router
      Next hop: 10.10.10.1 via xe-0/0/1.0
      Session Id: 0x2
    1.2.3.4/32 Originating RIB: inet.0
      Metric: 1
      Forwarding nexthops: 2
      Nexthop: 10.92.78.102 via em0.0
      Node path count: 1

3:1000:10::300::1.2.3.4/304 (1 entry, 0 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 1000:10
    PMSI: Flags 0x0: Label 300: Type INGRESS-REPLICATION 1.2.3.4
    Next hop type: Indirect
    Address: 0x9420fd0
    Next-hop reference count: 12
    Source: 1.2.3.4
    Protocol next hop: 1.2.3.4
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    State: Local AS: 17 Peer AS: 17 Age 35:22 Metric2:1 Validation State:
    unverified Task: BGP 17.1.2.3.4+5075
    6 AS path: I Communities: target:3333:33 encapsulation0:0:0:0:3
Import Accepted Localpref:100
  Router ID: 1.2.3.4
  Secondary Tables: default-switch.evpn.0
  Indirect next hops: 1
    Protocol next hop: 1.2.3.4 Metric: 1
    Indirect next hop: 0x2 no-forward INH Session ID: 0x0
    Indirect path forwarding next hops: 1
      Next hop type: Router
      Next hop: 10.10.10.1 via xe-0/0/1.0
      Session Id: 0x2
    1.2.3.4/32 Originating RIB: inet.0
      Metric: 1
      Forwarding nexthops: 2
      Nexthop: 10.92.78.102 via em0.0
      Node path count: 1

```

## show route forwarding-table

<b>List of Syntax</b>	<a href="#">Syntax on page 437</a> <a href="#">Syntax (MX Series Routers) on page 437</a> <a href="#">Syntax (TX Matrix and TX Matrix Plus Routers) on page 437</a>
<b>Syntax</b>	<pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;all&gt; &lt;ccc interface-name&gt; &lt;destination destination-prefix&gt; &lt;family family   matching matching&gt; &lt;interface-name interface-name&gt; &lt;label name&gt; &lt;matching matching&gt; &lt;multicast&gt; &lt;table (default   logical-system-name/routing-instance-name   routing-instance-name)&gt; &lt;vlan (all   vlan-name)&gt; &lt;vpn vpn&gt;</pre>
<b>Syntax (MX Series Routers)</b>	<pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;all&gt; &lt;bridge-domain (all   domain-name)&gt; &lt;ccc interface-name&gt; &lt;destination destination-prefix&gt; &lt;family family   matching matching&gt; &lt;interface-name interface-name&gt; &lt;label name&gt; &lt;learning-vlan-id learning-vlan-id&gt; &lt;matching matching&gt; &lt;multicast&gt; &lt;table (default   logical-system-name/routing-instance-name   routing-instance-name)&gt; &lt;vlan (all   vlan-name)&gt; &lt;vpn vpn&gt;</pre>
<b>Syntax (TX Matrix and TX Matrix Plus Routers)</b>	<pre>show route forwarding-table &lt;detail   extensive   summary&gt; &lt;all&gt; &lt;ccc interface-name&gt; &lt;destination destination-prefix&gt; &lt;family family   matching matching&gt; &lt;interface-name interface-name&gt; &lt;matching matching&gt; &lt;label name&gt; &lt;lcc number&gt; &lt;multicast&gt; &lt;table routing-instance-name&gt; &lt;vpn vpn&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Option <b>bridge-domain</b> introduced in Junos OS Release 7.5</p> <p>Option <b>learning-vlan-id</b> introduced in Junos OS Release 8.4</p>

Options **all** and **vlan** introduced in Junos OS Release 9.6.

Command introduced in Junos OS Release 11.3 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

**Description** Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.



**NOTE:** The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

---

**Options** **none**—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

**detail | extensive | summary**—(Optional) Display the specified level of output.

**all**—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

**bridge-domain (all | bridge-domain-name)**—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

**ccc interface-name**—(Optional) Display route entries for the specified circuit cross-connect interface.

**destination destination-prefix**—(Optional) Destination prefix.

**family family**—(Optional) Display routing table entries for the specified family: **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

**interface-name interface-name**—(Optional) Display routing table entries for the specified interface.

**label name**—(Optional) Display route entries for the specified label.

**lcc number**—(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers, display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

**learning-vlan-id** *learning-vlan-id*—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

**matching** *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

**multicast**—(Optional) Display routing table entries for multicast routes.

**table** (**default** | *logical-system-name/routing-instance-name* | *routing-instance-name*)—(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the **show route instance** command.

**vlan** (**all** | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

**vpn** *vpn*—(Optional) Display routing table entries for a specified VPN.

**Required Privilege Level**

view

**List of Sample Output**

[show route forwarding-table on page 442](#)  
[show route forwarding-table detail on page 443](#)  
[show route forwarding-table destination extensive \(Weights and Balances\) on page 443](#)  
[show route forwarding-table extensive on page 444](#)  
[show route forwarding-table extensive \(RPF\) on page 445](#)  
[show route forwarding-table family mpls on page 446](#)  
[show route forwarding-table family vpls on page 446](#)  
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled\) on page 446](#)  
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled with MAC Statistics\) on page 447](#)  
[show route forwarding-table family vpls extensive on page 447](#)  
[show route forwarding-table table default on page 448](#)  
[show route forwarding-table table logical-system-name/routing-instance-name on page 449](#)

[show route forwarding-table vpn on page 450](#)

**Output Fields** [Table 37 on page 357](#) lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

**Table 53: show route forwarding-table Output Fields**

Field Name	Field Description	Level of Output
Logical system	Name of the logical system. This field is displayed if you specify the <b>table logical-system-name/routing-instance-name</b> option on a device that is configured for and supports logical systems.	All levels
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels
Address family	Address family (for example, IP, IPv6, ISO, MPLS, and VPLS).	All levels
Destination	Destination of the route.	<b>detail extensive</b>
Route Type (Type)	How the route was placed into the forwarding table. When the <b>detail</b> keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> <li><b>cloned (clon)</b>—(TCP or multicast only) Cloned route.</li> <li><b>destination (dest)</b>—Remote addresses directly reachable through an interface.</li> <li><b>destination down (iddn)</b>—Destination route for which the interface is unreachable.</li> <li><b>interface cloned (ifcl)</b>—Cloned route for which the interface is unreachable.</li> <li><b>route down (ifdn)</b>—Interface route for which the interface is unreachable.</li> <li><b>ignore (ignr)</b>—Ignore this route.</li> <li><b>interface (intf)</b>—Installed as a result of configuring an interface.</li> <li><b>permanent (perm)</b>—Routes installed by the kernel when the routing table is initialized.</li> <li><b>user</b>—Routes installed by the routing protocol process or as a result of the configuration.</li> </ul>	All levels
Route Reference (RtRef)	Number of routes to reference.	<b>detail extensive</b>
Flags	Route type flags: <ul style="list-style-type: none"> <li><b>none</b>—No flags are enabled.</li> <li><b>accounting</b>—Route has accounting enabled.</li> <li><b>cached</b>—Cache route.</li> <li><b>incoming-iface interface-number</b>—Check against incoming interface.</li> <li><b>prefix load balance</b>—Load balancing is enabled for this prefix.</li> <li><b>rt nh decoupled</b>—Route has been decoupled from the next hop to the destination.</li> <li><b>sent to PFE</b>—Route has been sent to the Packet Forwarding Engine.</li> <li><b>static</b>—Static route.</li> </ul>	<b>extensive</b>
Next hop	IP address of the next hop to the destination.	<b>detail extensive</b>



Table 53: show route forwarding-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
Next hop Type (Type)	<p>Next-hop type. When the <b>detail</b> keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> <li>• <b>broadcast (bcst)</b>—Broadcast.</li> <li>• <b>deny</b>—Deny.</li> <li>• <b>discard (dscd)</b>—Discard.</li> <li>• <b>hold</b>—Next hop is waiting to be resolved into a unicast or multicast type.</li> <li>• <b>indexed (idxd)</b>—Indexed next hop.</li> <li>• <b>indirect (indr)</b>—Indirect next hop.</li> <li>• <b>local (locl)</b>—Local address on an interface.</li> <li>• <b>routed multicast (mcrst)</b>—Regular multicast next hop.</li> <li>• <b>multicast (mcst)</b>—Wire multicast next hop (limited to the LAN).</li> <li>• <b>multicast discard (mdsc)</b>—Multicast discard.</li> <li>• <b>multicast group (mgrp)</b>—Multicast group member.</li> <li>• <b>receive (rcv)</b>—Receive.</li> <li>• <b>reject (rjct)</b>—Discard. An ICMP unreachable message was sent.</li> <li>• <b>resolve (rslv)</b>—Resolving the next hop.</li> <li>• <b>unicast (ucst)</b>—Unicast.</li> <li>• <b>unilist (ulst)</b>—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.</li> </ul>	<b>detail extensive</b>
Index	Software index of the next hop that is used to route the traffic for a given prefix.	<b>detail extensive none</b>
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	<b>extensive</b>
Reference (NhRef)	Number of routes that refer to this next hop.	<b>detail extensive none</b>
Next-hop interface (Netif)	Interface used to reach the next hop.	<b>detail extensive none</b>
Weight	Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the <b>Balance</b> field description).	<b>extensive</b>
Balance	Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.	<b>extensive</b>
RPF interface	List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when <b>rpf-check</b> is configured on the interface.	<b>extensive</b>

## Sample Output

### show route forwarding-table

```

user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  46   4
0.0.0.0/32       perm  0                               dscd  44   1
1.1.1.0/24       ifdn  0                               rslv  608  1 ge-2/0/1.0
1.1.1.0/32       iddn  0 1.1.1.0             recv  606  1 ge-2/0/1.0
1.1.1.1/32       user  0                               rjct  46   4
1.1.1.1/32       intf  0 1.1.1.1             locl  607  2
1.1.1.1/32       iddn  0 1.1.1.1             locl  607  2
1.1.1.255/32     iddn  0 ff:ff:ff:ff:ff:ff   bcst  605  1 ge-2/0/1.0
10.0.0.0/24      intf  0                               rslv  616  1 ge-2/0/0.0
10.0.0.0/32      dest  0 10.0.0.0            recv  614  1 ge-2/0/0.0
10.0.0.1/32      intf  0 10.0.0.1            locl  615  2
10.0.0.1/32      dest  0 10.0.0.1            locl  615  2
10.0.0.255/32    dest  0 10.0.0.255          bcst  613  1 ge-2/0/0.0
10.1.1.0/24      ifdn  0                               rslv  612  1 ge-2/0/1.0
10.1.1.0/32      iddn  0 10.1.1.0            recv  610  1 ge-2/0/1.0
10.1.1.1/32      user  0                               rjct  46   4
10.1.1.1/32      intf  0 10.1.1.1            locl  611  2
10.1.1.1/32      iddn  0 10.1.1.1            locl  611  2
10.1.1.255/32    iddn  0 ff:ff:ff:ff:ff:ff   bcst  609  1 ge-2/0/1.0
10.206.0.0/16    user  0 10.209.63.254        ucst  419  20 fxp0.0
10.209.0.0/16    user  1 0:12:1e:ca:98:0      ucst  419  20 fxp0.0
10.209.0.0/18    intf  0                               rslv  418  1 fxp0.0
10.209.0.0/32    dest  0 10.209.0.0          recv  416  1 fxp0.0
10.209.2.131/32  intf  0 10.209.2.131        locl  417  2
10.209.2.131/32  dest  0 10.209.2.131        locl  417  2
10.209.17.55/32  dest  0 0:30:48:5b:78:d2     ucst  435  1 fxp0.0
10.209.63.42/32  dest  0 0:23:7d:58:92:ca     ucst  434  1 fxp0.0
10.209.63.254/32 dest  0 0:12:1e:ca:98:0      ucst  419  20 fxp0.0
10.209.63.255/32 dest  0 10.209.63.255       bcst  415  1 fxp0.0
10.227.0.0/16    user  0 10.209.63.254        ucst  419  20 fxp0.0

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  27   1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf  0                               locl  28   1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  6    1
ff00::/8         perm  0                               mdsc  4    1
ff02::1/128      perm  0 ff02::1             mcst  3    1

Routing table: ccc
MPLS:
Interface.Label  Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                               rjct  16   1
100004(top)fe-0/0/1.0

```

## show route forwarding-table detail

```

user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          user   2 0:90:69:8e:b1:1b ucst  132   4 fxp0.0
default          perm   0                               rjct  14    1
10.1.1.0/24      intf   0 ff.3.0.21         ucst  322   1 so-5/3/0.0
10.1.1.0/32      dest   0 10.1.1.0          recv  324   1 so-5/3/0.0
10.1.1.1/32      intf   0 10.1.1.1          locl  321   1
10.1.1.255/32    dest   0 10.1.1.255        bcst  323   1 so-5/3/0.0
10.21.21.0/24    intf   0 ff.3.0.21         ucst  326   1 so-5/3/0.0
10.21.21.0/32    dest   0 10.21.21.0        recv  328   1 so-5/3/0.0
10.21.21.1/32    intf   0 10.21.21.1        locl  325   1
10.21.21.255/32  dest   0 10.21.21.255      bcst  327   1 so-5/3/0.0
127.0.0.1/32     intf   0 127.0.0.1         locl  320   1
172.17.28.19/32  clon   1 192.168.4.254     ucst  132   4 fxp0.0
172.17.28.44/32  clon   1 192.168.4.254     ucst  132   4 fxp0.0

...

Routing table: private1__inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct  46    1
10.0.0.0/8       intf   0                               rslv  136   1 fxp1.0
10.0.0.0/32      dest   0 10.0.0.0          recv  134   1 fxp1.0
10.0.0.4/32      intf   0 10.0.0.4          locl  135   2
10.0.0.4/32      dest   0 10.0.0.4          locl  135   2

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct  38    1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct  22    1
ff00::/8         perm   0                               mdsc  21    1
ff02::1/128      perm   0 ff02::1          mcst  17    1

...

Routing table: mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct  28    1

```

## show route forwarding-table destination extensive (Weights and Balances)

```

user@host> show route forwarding-table destination 3.4.2.1 extensive
Routing table: inet [Index 0]
Internet:

Destination: 3.4.2.1/32
Route type: user
Route reference: 0                               Route interface-index: 0

```

Flags: sent to PFE		
Next-hop type: unicast	Index: 262143	Reference: 1
Nexthop: 4.4.4.4		
Next-hop type: unicast	Index: 335	Reference: 2
Next-hop interface: so-1/1/0.0	Weight: 22	Balance: 3
Nexthop: 145.12.1.2		
Next-hop type: unicast	Index: 337	Reference: 2
Next-hop interface: so-0/1/2.0	Weight: 33	Balance: 33

### show route forwarding-table extensive

```
user@host> show route forwarding-table extensive
```

```
Routing table: inet [Index 0]
```

```
Internet:
```

```
Destination: default
```

```
Route type: user
```

```
Route reference: 2
```

```
Route interface-index: 0
```

```
Flags: sent to PFE
```

```
Nexthop: 0:90:69:8e:b1:1b
```

```
Next-hop type: unicast
```

```
Index: 132      Reference: 4
```

```
Next-hop interface: fxp0.0
```

```
Destination: default
```

```
Route type: permanent
```

```
Route reference: 0
```

```
Route interface-index: 0
```

```
Flags: none
```

```
Next-hop type: reject
```

```
Index: 14      Reference: 1
```

```
Destination: 127.0.0.1/32
```

```
Route type: interface
```

```
Route reference: 0
```

```
Route interface-index: 0
```

```
Flags: sent to PFE
```

```
Nexthop: 127.0.0.1
```

```
Next-hop type: local
```

```
Index: 320      Reference: 1
```

```
...
```

```
Routing table: private1__inet [Index 1]
```

```
Internet:
```

```
Destination: default
```

```
Route type: permanent
```

```
Route reference: 0
```

```
Route interface-index: 0
```

```
Flags: sent to PFE
```

```
Next-hop type: reject
```

```
Index: 46      Reference: 1
```

```
Destination: 10.0.0.0/8
```

```
Route type: interface
```

```
Route reference: 0
```

```
Route interface-index: 3
```

```
Flags: sent to PFE
```

```
Next-hop type: resolve
```

```
Index: 136      Reference: 1
```

```
Next-hop interface: fxp1.0
```

```
...
```

```
Routing table: iso [Index 0]
```

```
ISO:
```

```
Destination: default
```

```
Route type: permanent
```

```

Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 38      Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 22      Reference: 1

Destination: ff00::/8
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: multicast discard
Route interface-index: 0
Index: 21      Reference: 1

...

Routing table: private1__inet6 [Index 1]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 54      Reference: 1

Destination: fe80::2a0:a5ff:fe3d:375/128
Route type: interface
Route reference: 0
Flags: sent to PFE
Next-hop: fe80::2a0:a5ff:fe3d:375
Next-hop type: local
Route interface-index: 0
Index: 75      Reference: 1

...

```

### show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```

so-1/1/0 {
  unit 0 {
    family inet {
      rpf-check;
      address 15.95.1.2/30;
    }
  }
}

```

```

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:
...
...
Destination: 15.95.1.3/32
Route type: destination
Route reference: 0
Route interface-index: 67

```

```

Flags: sent to PFE
Nexthop: 15.95.1.3
Next-hop type: broadcast          Index: 328      Reference: 1
Next-hop interface: so-1/1/0.0
RPF interface: so-1/1/0.0

```

### show route forwarding-table family mpls

```

user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0
0                user  0
1                user  0
2                user  0
100000           user  0 10.31.1.6      swap 100001      fe-1/1/0.0
800002           user  0                Pop          vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
                user  0                indr  351      4
                Push 800000, Push 100002(top)

so-0/0/0.0

```

### show route forwarding-table family vpls

```

user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          dymn  0
default          perm  0
fe-0/1/0.0       dymn  0
00:90:69:0c:20:1f/48      <<<<<Remote CE

                dymn  0                indr  351      4
                Push 800000, Push 100002(top)

so-0/0/0.0
00:90:69:85:b0:1f/48      <<<<<Local CE

                dymn  0                ucst  354      2 fe-0/1/0.0

```

### show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled)

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing
Destination      Type RtRef Next hop      Type Index      NhRef Netif
default          perm  0
lsi.1048832      intf  0
                4.4.3.2      indr 1048574      4
                Push 262145      621      2

ge-3/0/0.0
00:19:e2:25:d0:01/48 user  0                ucst  590      5 ge-2/3/9.0
0x30003/51       user  0                comp  627      2
ge-2/3/9.0       intf  0                ucst  590      5 ge-2/3/9.0
ge-3/1/3.0       intf  0                ucst  619      4 ge-3/1/3.0
0x30002/51       user  0                comp  600      2
0x30001/51       user  0                comp  597      2

```

### show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics)

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats
Destination      Type RtRef Next hop      Type Index  NhRef Netif
default          perm  0         4.4.3.2      dscd   519      1
1si.1048834      intf  0         4.4.3.2      indr  1048574  4
                  Push  262145    592      2
ge-3/0/0.0
00:19:e2:25:d0:01/48 user  0         ucst   590      5 ge-2/3/9.0
0x30003/51      user  0         comp   630      2
ge-2/3/9.0      intf  0         ucst   590      5 ge-2/3/9.0
ge-3/1/3.0      intf  0         ucst   591      4 ge-3/1/3.0
0x30002/51      user  0         comp   627      2
0x30001/51      user  0         comp   624      2

```

### show route forwarding-table family vpls extensive

```

user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

Destination: default
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: unicast
Next-hop interface: fe-0/1/3.0
Next-hop type: unicast
Next-hop interface: fe-0/1/2.0
Route interface-index: 72
Index: 289      Reference: 1
Index: 291      Reference: 3
Index: 290      Reference: 3

Destination: default
Route type: permanent
Route reference: 0
Flags: none
Next-hop type: discard
Route interface-index: 0
Index: 341      Reference: 1

Destination: fe-0/1/2.0
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: indirect
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast
Next-hop interface: fe-0/1/3.0
Route interface-index: 69
Index: 293      Reference: 1
Index: 363      Reference: 4
Index: 301      Reference: 5
Index: 291      Reference: 3

Destination: fe-0/1/3.0
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Route interface-index: 70
Index: 292      Reference: 1

```

```

Next-hop type: indirect          Index: 363      Reference: 4
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect          Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast           Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0

Destination: 10:00:00:01:01:01/48
Route type: dynamic
Route reference: 0                Route interface-index: 70
Flags: sent to PFE, prefix load balance
Next-hop type: unicast           Index: 291      Reference: 3
Next-hop interface: fe-0/1/3.0
Route used as destination:
  Packet count:      6640    Byte count:      675786
Route used as source
  Packet count:      6894    Byte count:      696424

Destination: 10:00:00:01:01:04/48
Route type: dynamic
Route reference: 0                Route interface-index: 69
Flags: sent to PFE, prefix load balance
Next-hop type: unicast           Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0
Route used as destination:
  Packet count:      96      Byte count:      8079
Route used as source:
  Packet count:      296     Byte count:      24955

Destination: 10:00:00:01:03:05/48
Route type: dynamic
Route reference: 0                Route interface-index: 74
Flags: sent to PFE, prefix load balance
Next-hop type: indirect          Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0

```

### show route forwarding-table table default

```

user@host> show route forwarding-table table default
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
0.0.0.0/32       perm  0
10.0.60.0/30     user  0 10.0.60.13             ucst  713  5 fe-0/1/3.0
10.0.60.12/30    intf  0                      rslv  688  1 fe-0/1/3.0
10.0.60.12/32    dest  0 10.0.60.12             recv  686  1 fe-0/1/3.0
10.0.60.13/32    dest  0 0:5:85:8b:bc:22        ucst  713  5 fe-0/1/3.0
10.0.60.14/32    intf  0 10.0.60.14             locl  687  2
10.0.60.14/32    dest  0 10.0.60.14             locl  687  2
10.0.60.15/32    dest  0 10.0.60.15             bcst  685  1 fe-0/1/3.0
10.0.67.12/30    user  0 10.0.60.13             ucst  713  5 fe-0/1/3.0
10.0.80.0/30     ifdn  0 ff.3.0.21             ucst  676  1 so-0/0/1.0
10.0.80.0/32     dest  0 10.0.80.0             recv  678  1 so-0/0/1.0
10.0.80.2/32     user  0                      rjct  36   2
10.0.80.2/32     intf  0 10.0.80.2             locl  675  1

```



```

10.0.80.3/32      dest    0 10.0.80.3      bcst   677    1 so-0/0/1.0
10.0.90.12/30     intf    0                rslv   684    1 fe-0/1/0.0
10.0.90.12/32     dest    0 10.0.90.12    recv   682    1 fe-0/1/0.0
10.0.90.14/32     intf    0 10.0.90.14     locl   683    2
10.0.90.14/32     dest    0 10.0.90.14     locl   683    2
10.0.90.15/32     dest    0 10.0.90.15     bcst   681    1 fe-0/1/0.0
10.5.0.0/16       user    0 192.168.187.126 ucst   324   15 fxp0.0
10.10.0.0/16      user    0 192.168.187.126 ucst   324   15 fxp0.0
10.13.10.0/23     user    0 192.168.187.126 ucst   324   15 fxp0.0
10.84.0.0/16      user    0 192.168.187.126 ucst   324   15 fxp0.0
10.150.0.0/16     user    0 192.168.187.126 ucst   324   15 fxp0.0
10.157.64.0/19    user    0 192.168.187.126 ucst   324   15 fxp0.0
10.209.0.0/16     user    0 192.168.187.126 ucst   324   15 fxp0.0

```

...

Routing table: default.iso

ISO:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	60	1	

Routing table: default.inet6

Internet6:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	44	1	
::/128	perm	0		dscd	42	1	
ff00::/8	perm	0		mdsc	43	1	
ff02::1/128	perm	0	ff02::1	mcst	39	1	

Routing table: default.mpls

MPLS:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		dscd	50	1	

### show route forwarding-table table logical-system-name/routing-instance-name

```
user@host> show route forwarding-table table R4/vpn-red
```

Logical system: R4

Routing table: vpn-red.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	563	1	
0.0.0.0/32	perm	0		dscd	561	2	
1.0.0.1/32	user	0		dscd	561	2	
2.0.2.0/24	intf	0		rslv	771	1	ge-1/2/0.3
2.0.2.0/32	dest	0	2.0.2.0	recv	769	1	ge-1/2/0.3
2.0.2.1/32	intf	0	2.0.2.1	locl	770	2	
2.0.2.1/32	dest	0	2.0.2.1	locl	770	2	
2.0.2.2/32	dest	0	0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0	ucst	789	1	ge-1/2/0.3
2.0.2.255/32	dest	0	2.0.2.255	bcst	768	1	ge-1/2/0.3
224.0.0.0/4	perm	1		mdsc	562	1	
224.0.0.1/32	perm	0	224.0.0.1	mcst	558	1	
255.255.255.255/32	perm	0		bcst	559	1	

Logical system: R4

Routing table: vpn-red.iso

ISO:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	608	1	

```

Logical system: R4
Routing table: vpn-red.inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct  708   1
::/128           perm  0                dscd  706   1
ff00::/8         perm  0                mdsc  707   1
ff02::1/128     perm  0 ff02::1          mcst  704   1

```

```

Logical system: R4
Routing table: vpn-red.mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                dscd  638

```

### show route forwarding-table vpn

```

user@host> show route forwarding-table vpn VPN-A
Routing table:: VPN-A.inet
Internet:
Destination      Type RtRef Nexthop          Type Index NhRef Netif
default          perm  0                rjct   4    4
10.39.10.20/30   intf  0 ff.3.0.21          ucst   40    1
so-0/0/0.0
10.39.10.21/32   intf  0 10.39.10.21        locl   36    1
10.255.14.172/32 user   0                ucst   69    2
so-0/0/0.0
10.255.14.175/32 user   0                indr   81    3
Push 100004, Push
100004(top) so-1/0/0.0
224.0.0.0/4      perm  2                mdsc   5    3
224.0.0.1/32     perm  0 224.0.0.1          mcst   1    8
224.0.0.5/32     user   1 224.0.0.5          mcst   1    8
255.255.255.255/32 perm  0                bcst   2    3

```

## show mpls static-lsp

**Syntax** show mpls static-lsp  
 <brief | detail | extensive | terse>  
 <bypass>  
 <descriptions>  
 <down | up>  
 <ingress>  
 <instance *instance-name*>  
 <logical-system (all | *logical-system-name*)>  
 <lsp-type>  
 <name *name*>  
 <statistics>  
 <transit>

**Release Information** Command introduced in Junos OS Release 10.1.  
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Display information about configured and active static Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

**Options** **none**—Display standard information about all configured and active static MPLS LSPs.

**brief | detail | extensive | terse**—(Optional) Display the specified level of output. The **extensive** option displays the same information as the **detail** option, but covers the most recent 50 events.

**bypass**—(Optional) Display LSPs used for protecting other static LSPs.

**descriptions**—(Optional) Display the MPLS static LSP descriptions. To view this information, you must configure the description statement at the **[edit protocols mpls static-label-switched-path *path-name* bypass]**, **[edit protocols mpls static-label-switched-path *path-name* ingress]**, or **[edit protocols mpls static-label-switched-path *path-name* transit *incoming-label*]** hierarchy levels. Only static LSPs with a description are displayed.

**down | up**—(Optional) Display only static LSPs that are inactive or active, respectively.

**instance *instance-name***—(Optional) Display information about all configured and active static MPLS LSPs for the specified routing instance. If ***instance-name*** is omitted, information about all configured and active static MPLS LSPs for the master instance is displayed.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

***lsp-type***—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

**name *name***—(Optional) Display information about the specified static LSP or group of LSPs.

**statistics**—(Optional) Display accounting information about static LSPs.

**transit**—(Optional) Display static LSPs transiting this routing device.

**Required Privilege Level** view

**List of Sample Output** [show mpls static-lsp extensive on page 453](#)  
[show mpls static-lsp statistics ingress on page 453](#)  
[show mpls static-lsp \(when MPLS stitching is used\) on page 453](#)

**Output Fields** [Table 54 on page 452](#) describes the output fields for the **show mpls static-lsp** command. Output fields are listed in the approximate order in which they appear.

**Table 54: show mpls static-lsp Output Fields**

Field Name	Field Description	Level of Output
<b>Ingress LSPs</b>	Information about the static LSPs on the ingress routing device. Each session has one line of output.	All levels
<b>Transit LSPs</b>	Number of static LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
<b>Bypass LSPs</b>	Information about the bypass LSPs configured on the routing device. Each session has one line of output.	All levels
<b>LSPname</b>	Name of the static LSP.	All levels
<b>To</b>	Destination (egress routing device) of the session.	All levels
<b>State</b>	State of the static LSP handled by this RSVP session: <b>Up</b> , <b>Dn</b> (down), or <b>Restart</b> .	All levels
<b>Packets</b>	Number of packet transiting the static LSP ( <b>statistics</b> option only).	All levels
<b>Bytes</b>	Number of bytes transiting the static LSP ( <b>statistics</b> option only).	All levels
<b>Nexthop</b>	IP address for the next-hop router for the static LSP.	<b>detail, extensive</b>
<b>Bypass</b>	(Bypass LSP) Destination address (egress routing device) for the bypass LSP.	All levels
<b>Link protection desired</b>	Link protection has been requested by the ingress routing device.	<b>detail, extensive</b>
<b>LabelOperation</b>	Label operation to perform: <b>Push</b> , <b>Pop</b> , <b>Swap</b> .	<b>detail, extensive</b>
<b>Outgoing-label</b>	Outgoing label to use for the MPLS packet in either push or swap label operations.	<b>detail, extensive</b>

Table 54: show mpls static-lsp Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Created</b>	(Ingress LSP) Date and time the static LSP was created.	<b>extensive</b>
<b>Bandwidth</b>	Bandwidth configured for the static LSP.	<b>detail, extensive</b>
<b>Resv style</b>	(Bypass) RSVP reservation style. This field consists of two parts: the number of active reservations and the reservation style, which can be <b>FF</b> (fixed filter), <b>SE</b> (shared explicit), or <b>WF</b> (wildcard filter).	All levels

## Sample Output

### show mpls static-lsp extensive

```

user@host> show mpls static-lsp extensive
Ingress LSPs:
LSPName: alpha-to-beta, To: 192.168.14.1
State: Dn
Nexthop: 192.168.10.1
LabelOperation: Push, Outgoing-label: 1000001
Created: Thu Jan 14 16:44:43 2010
Bandwidth: 0 bps
Total 1, displayed 1, Up 0, Down 1

Transit LSPs:
Total 0, displayed 0, Up 0, Down 0

Bypass LSPs:
Total 0, displayed 0, Up 0, Down 0

```

### show mpls static-lsp statistics ingress

```

user@host> show mpls static-lsp statistics ingress
Ingress LSPs:
LSPName                To           State    Packets    Bytes
alpha-to-beta          192.168.14.1 Dn        NA         NA
Total 1, displayed 1, Up 0, Down 1

```

### show mpls static-lsp (when MPLS stitching is used)

The show mpls static-lsp command was extended in Junos release 14.1X53-D25 to accommodate the stitching feature of MPLS. This example shows the LSP state as 'InProgress' because the LSP is waiting for protocol next-hop resolution. For more information, see

```

user@host> show mpls static-lsp
Ingress LSPs:
Total 0, displayed 0, Up 0, Down 0
Transit LSPs: LSPName                Incoming-label  State
to-165        1000000        InProgress

```

## show ted database

<b>List of Syntax</b>	<a href="#">Syntax on page 454</a> <a href="#">Syntax (EX Series Switches) on page 454</a>
<b>Syntax</b>	<pre>show ted database &lt;brief   detail   extensive&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;<i>system-name</i>&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show ted database &lt;brief   detail   extensive&gt; &lt;<i>system-name</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display the entries in the Multiprotocol Label Switching (MPLS) traffic engineering database.
<b>Options</b>	<p><b>none</b>—Display standard information about all entries in the traffic engineering database.</p> <p><b>brief   detail   extensive</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display routing instance information for the specified instance. If <i>instance-name</i> is omitted, information is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b><i>system-name</i></b>—(Optional) Display traffic engineering database information for a particular system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ted database brief on page 457</a> <a href="#">show ted database detail on page 457</a> <a href="#">show ted database extensive on page 458</a>
<b>Output Fields</b>	<p><a href="#">Table 55 on page 454</a> describes the output fields for the <b>show ted database</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 55: show ted database Output Fields**

Field Name	Field Description	Level of Output
TED database	Number of nodes and pseudonodes participating in IS-IS and OSPF domain routing.	All levels

Table 55: show ted database Output Fields (*continued*)

Field Name	Field Description	Level of Output
ID	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode. If the node contains a router ID, it is displayed in parentheses.	<b>brief</b>
NodeID	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	<b>extensive</b>
Type	Type of node. It can be either <b>Rtr</b> (router) or <b>Net</b> (pseudonode).	All levels
Age(s)	How long since the node was last refreshed, in seconds.	All levels
LnkIn	Number of nodes pointing toward this node.	All levels
LnkOut	Number of nodes to which this node points.	All levels
Protocol	Protocol that reported the node information: <ul style="list-style-type: none"> <li>• <b>IS-IS(1)</b>—IS-IS Level 1.</li> <li>• <b>IS-IS(2)</b>—IS-IS Level 2.</li> <li>• <b>OSPF (area-number)</b>—OSPF from the specified area.</li> </ul>	All levels
To	Address on the far end of a link.	<b>detail extensive</b>
Local	Address of the local interface being used to reach the remote node.	<b>detail extensive</b>
Remote	Address of the interface on the remote node.	<b>detail extensive</b>
Local interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	<b>detail extensive</b>
Remote interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	<b>detail extensive</b>
Metric	Configured traffic engineering metric.	<b>extensive</b>
IGP metric	Configured interior gateway protocol metric.	<b>extensive</b>
Static BW	Total interface bandwidth in bps.	<b>extensive</b>
Reservable bandwidth	Subscription factor for the interface, which is the percentage of the link bandwidth that can be used for the RSVP reservation process. You configure this by including the <b>subscription</b> statement when configuring RSVP.	<b>extensive</b>
<b>Available BW [priority]</b>	(Must include <b>diffserv-te</b> statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each priority level. The bandwidth shown is for the entire interface, not for each individual LSP.	<b>extensive</b>

Table 55: show ted database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Diffserv-TE BW Model	Bandwidth constraint model used by the LSPs.	extensive
Available BW [TE-class]	(Must include the <b>diffserv-te</b> statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each traffic engineering class.	extensive
Static BW [CT-class]	Total interface bandwidth used by an MPLS traffic class, in bps.	extensive
Interface Switching Capability Descriptor ( <i>n</i> )	<p>Information about the interface switching capability descriptor, which is a subtype length value (TLV) of the link TLV. <i>n</i> is the index number.</p> <ul style="list-style-type: none"> <li>• <b>Switching type</b>—Type of switching to be performed on a particular link: <ul style="list-style-type: none"> <li>• PSC-1—Packet switch-capable 1</li> <li>• PSC-2—Packet switch-capable 2</li> <li>• PSC-3—Packet switch-capable 3</li> <li>• PSC-4—Packet switch-capable 4</li> <li>• L2SC—Layer-2-switch-capable</li> <li>• TDM—Time-division-multiplexing-capable</li> <li>• LSC—Lambda switch-capable</li> <li>• FSC—Fiber switch-capable</li> </ul> </li> <li>• <b>Encoding type</b>—Encoding of the LSP being requested: <ul style="list-style-type: none"> <li>• Packet</li> <li>• Ethernet</li> <li>• ANSI/ETSI PDH</li> <li>• Reserved</li> <li>• SDH /SONET</li> <li>• Digital Wrapper</li> <li>• Lambda (photonic)</li> <li>• Fiber</li> <li>• FiberSDH/SONET</li> </ul> </li> <li>• <b>Maximum LSP BW [priority] bps</b>—Maximum LSP bandwidth information. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. <ul style="list-style-type: none"> <li>• [<i>n</i>]—Priority level. The range is from 0 (high) through 7 (low).</li> <li>• <i>n</i> Mbps—Amount of the maximum bandwidth.</li> </ul> </li> <li>• <b>Minimum LSP BW</b>—Minimum LSP bandwidth in Mbps. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. <b>Minimum LSP BW</b> is displayed only when <b>switching type</b> is PSC-1 or TDM.</li> <li>• <b>Interface MTU</b>—Displayed only when <b>switching type</b> is TDM.</li> <li>• <b>Interface supports standard SONET/SDH</b>—Displayed only when <b>switching type</b> is TDM.</li> </ul>	extensive



## Sample Output

### show ted database brief

```

user@host> show ted database brief
TED database: 12 ISIS nodes 0 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-A.00                      ---  3178    2    0
Router-B.00                      ---  3152    2    0
Router-B.02                      Net   802    0    2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-C.00                      ---  3126    2    0
Router-C.02                      Net   38     0    2 IS-IS(2)
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-D.00                      ---  3144    2    0
Router-D.02                      Net   723    0    2 IS-IS(2)
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-D.03                      Net   607    0    2 IS-IS(2)
    To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-E.00                      ---  3178    2    0
Router-E.02                      Net   131    0    2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-F.00                      ---  3153    2    0
Router-F.02                      Net   769    0    2 IS-IS(2)
    To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0

```

### show ted database detail

```

TED database: 12 ISIS nodes 0 INET nodes
ID                               Type Age(s) LnkIn LnkOut Protocol
Router-A.00                      ---  2913    2    0
Router-B.00                      ---  2887    2    0
Router-B.02                      Net   537    0    2 IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                               Type Age(s) LnkIn LnkOut Protocol

```

```

Router-C.00          ---      2861      2      0
Router-C.02          Net       597      0      2 IS-IS(2)
  To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                   Type Age(s) LnkIn LnkOut Protocol
Router-D.00          ---      2879      2      0
Router-D.02          Net       458      0      2 IS-IS(2)
  To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                   Type Age(s) LnkIn LnkOut Protocol
Router-D.03          Net       342      0      2 IS-IS(2)
  To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                   Type Age(s) LnkIn LnkOut Protocol
Router-E.00          ---      2913      2      0
Router-E.02          Net       640      0      2 IS-IS(2)
  To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
ID                   Type Age(s) LnkIn LnkOut Protocol
Router-F.00          ---      2888      2      0
Router-F.02          Net       504      0      2 IS-IS(2)
  To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0
  To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
    Local interface index: 0, Remote interface index: 0

```

### show ted database extensive

```

user@host> show ted database extensive
TED database: 12 ISIS nodes 0 INET nodes
NodeID: Router-A.00
  Type: ---, Age: 3067 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-B.00
  Type: ---, Age: 3041 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-B.02
  Type: Net, Age: 691 secs, LinkIn: 0, LinkOut: 2
  Protocol: IS-IS(2)
    To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
      Metric: 0
      IGP metric: 10
      Interface Switching Capability Descriptor(1):
        Switching type: Packet
        Encoding type: Packet
        Maximum LSP BW [priority] bps:
          [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
          [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
    To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
      Local interface index: 0, Remote interface index: 0
      Metric: 0
      IGP metric: 20
      Interface Switching Capability Descriptor(1):
        Switching type: Packet

```

```

        Encoding type: Packet
        Maximum LSP BW [priority] bps:
            [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
            [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
NodeID: Router-C.00
Type: ---, Age: 3015 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-C.02
Type: Net, Age: 751 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10      Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
NodeID: Router-D.00
Type: ---, Age: 3034 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-D.02
Type: Net, Age: 613 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps      [1] 0bps      [2] 0bps      [3] 0bps
        [4] 0bps      [5] 0bps      [6] 0bps      [7] 0bps
NodeID: Router-D.03
Type: Net, Age: 497 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):

```

```

Switching type: Packet
Encoding type: Packet
Maximum LSP BW [priority] bps:
    [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
    [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
To: Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
        [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
NodeID: Router-E.00
Type: ---, Age: 3068 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-E.02
Type: Net, Age: 21 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
        [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
        [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
NodeID: Router-F.00
Type: ---, Age: 3043 secs, LinkIn: 2, LinkOut: 0
NodeID: Router-F.02
Type: Net, Age: 659 secs, LinkIn: 0, LinkOut: 2
Protocol: IS-IS(2)
To: Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:
        [0] 0bps    [1] 0bps    [2] 0bps    [3] 0bps
        [4] 0bps    [5] 0bps    [6] 0bps    [7] 0bps
To: Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
Local interface index: 0, Remote interface index: 0
Metric: 0
IGP metric: 10
Interface Switching Capability Descriptor(1):
    Switching type: Packet
    Encoding type: Packet
    Maximum LSP BW [priority] bps:

```

[0] 0bps	[1] 0bps	[2] 0bps	[3] 0bps
[4] 0bps	[5] 0bps	[6] 0bps	[7] 0bps

## show ted link

<b>List of Syntax</b>	<a href="#">Syntax on page 462</a> <a href="#">Syntax (EX Series Switches) on page 462</a>
<b>Syntax</b>	<pre>show ted link &lt;brief   detail&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show ted link &lt;brief   detail&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display Multiprotocol Label Switching (MPLS) traffic engineering database link information.
<b>Options</b>	<p><b>none</b>—Display standard information about traffic engineering database link information.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display routing instance information for the specified instance. If <b><i>instance-name</i></b> is omitted, information is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ted link brief on page 463</a> <a href="#">show ted link detail on page 463</a>
<b>Output Fields</b>	<a href="#">Table 56 on page 462</a> describes the output fields for the <b>show ted link</b> command. Output fields are listed in the approximate order in which they appear.

Table 56: show ted link Output Fields

Field Name	Field Description	Level of Output
ID	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	<b>brief</b>
-->ID	Hostname and address of the node that the link is going to. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	<b>brief</b>

Table 56: show ted link Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>hostname</i>	Hostname and address of the node that the link is coming from. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	<b>detail</b>
<i>hostname</i>	Hostname and address of the node that the link is going to. An address of .00 indicates that the node is the routing device itself. An address in the range 0.01 through 0.FF indicates that the node is a pseudonode.	<b>detail</b>
Local Path	Number of paths CSPF on the local routing device has placed on the link.	All levels
Metric	Configured traffic engineering metric.	<b>extensive</b>
IGP metric	Configured interior gateway protocol metric.	<b>detail</b>
Local BW	Amount of bandwidth the local routing device has placed on the link.	All levels
Local	Address of the local interface being used to reach the remote node.	<b>detail extensive</b>
Remote	Address of the interface on the remote node.	<b>detail extensive</b>
Local interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	<b>detail</b>
Remote interface index	The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.	<b>detail</b>

## Sample Output

### show ted link brief

```

user@host> show ted link brief
ID                               ->ID                               LocalPath LocalBW
Router-B.02                     Router-A.00                       0 0bps
Router-B.02                     Router-B.00                       0 0bps
Router-C.02                     Router-B.00                       0 0bps
Router-C.02                     Router-C.00                       0 0bps
Router-D.02                     Router-F.00                       0 0bps
Router-D.02                     Router-D.00                       0 0bps
Router-D.03                     Router-D.00                       0 0bps
Router-D.03                     Router-C.00                       0 0bps
Router-E.02                     Router-A.00                       0 0bps
Router-E.02                     Router-E.00                       0 0bps
Router-F.02                     Router-E.00                       0 0bps
Router-F.02                     Router-F.00                       0 0bps

```

### show ted link detail

```

user@host> show ted link detail
Router-B.02->Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: 0bps
  localBW [0] 0bps [1] 0bps [2] 0bps [3] 0bps

```

```
    localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-B.02->Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 20 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-C.02->Router-B.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 40 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-C.02->Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-D.02->Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-D.02->Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 60 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-D.03->Router-D.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-D.03->Router-C.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-E.02->Router-A.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 60 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-E.02->Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 20 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-F.02->Router-E.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 10 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
Router-F.02->Router-F.00, Local: 0.0.0.0, Remote: 0.0.0.0
  Local interface index: 0, Remote interface index: 0
  LocalPath: 0, Metric: 0, IGP metric: 40 AvailBW: Obps
  localBW [0] Obps [1] Obps [2] Obps [3] Obps
  localBW [4] Obps [5] Obps [6] Obps [7] Obps
```



## show ted protocol

<b>List of Syntax</b>	<a href="#">Syntax on page 465</a> <a href="#">Syntax (EX Series Switches) on page 465</a>
<b>Syntax</b>	<pre>show ted protocol &lt;brief   detail&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show ted protocol &lt;brief   detail&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display information about the protocols from which the Multiprotocol Label Switching (MPLS) traffic engineering database learned about its nodes.
<b>Options</b>	<p><b>none</b>—Display standard information about the protocols from which the traffic engineering database learned about its nodes.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display routing instance information for the specified instance. If <b><i>instance-name</i></b> is omitted, information is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ted protocol on page 466</a>
<b>Output Fields</b>	<p><a href="#">Table 57 on page 465</a> describes the output fields for the <b>show ted protocol</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 57: show ted protocol Output Fields**

Field Name	Field Description
<b>Protocol name</b>	<p>Protocol that reported the node information:</p> <ul style="list-style-type: none"> <li><b>IS-IS(1)</b>—IS-IS Level 1.</li> <li><b>IS-IS(2)</b>—IS-IS Level 2.</li> <li><b>OSPF (<i>area-number</i>)</b>—OSPF from the specified area.</li> </ul>
<b>Credibility</b>	If the protocols provide conflicting information about a node, the protocol with the highest credibility value is the one that the traffic engineering database uses.
<b>Self node</b>	Address the protocol uses as the local address.

## Sample Output

show ted protocol

```
user@host> show ted protocol
Protocol name      Credibility Self node
IS-IS(2)           2 (highest) corriedale.00(123.456.1.11)
IS-IS(1)           1          corriedale.00(123.456.1.11)
```

## PART 3

# RSVP

- [Using RSVP on page 469](#)
- [Configuration Statements for RSVP on page 511](#)
- [Monitoring Commands for RSVP on page 551](#)



## CHAPTER 7

# Using RSVP

- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 469](#)
- [RSVP Overview on page 473](#)
- [MTU Signaling in RSVP on page 474](#)
- [Tunneling LDP LSPs in RSVP LSPs on page 474](#)
- [Tunneling LDP LSPs in RSVP LSPs Overview on page 475](#)
- [Configuring MPLS on Provider Edge Switches on page 475](#)
- [Configuring MPLS on Provider Switches on page 479](#)
- [Verifying That MPLS Is Working Correctly on page 480](#)
- [Dynamic Bandwidth Management Using Container LSP Overview on page 482](#)

### Understanding MPLS Components for QFX Series and EX4600 Switches

---

MPLS devices include a number of components. While some components are required for all MPLS applications, others might not be, depending on the specific application.

This topic includes:

- [Provider Edge Switches on page 469](#)
- [Provider Switch on page 471](#)
- [Components Required for All Switches in the MPLS Network on page 471](#)

### Provider Edge Switches

To implement MPLS on a network, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch. In addition, you must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets.

The ingress PE switch (the entry point to the MPLS tunnel) receives a packet, analyzes it, and pushes an MPLS label onto it. This label places the packet in a forwarding equivalence class (FEC) and determines its handling and destination through the MPLS tunnel. The egress PE switch (the exit point from the MPLS tunnel) pops the MPLS label off the outgoing packet.

Within an MPLS tunnel, the network traffic is bidirectional. Therefore, each PE switch can be configured to be both an ingress switch and an egress switch, depending on the direction of the traffic.

The following MPLS components are configured on the PE switches but not on the provider switches:

- [MPLS Protocol and Label-Switched Paths on page 470](#)
- [IP Over MPLS for Customer Edge Interfaces on page 470](#)
- [BGP Layer 3 VPN Configuration on page 470](#)
- [Routing Instances for Layer 3 VPN on page 470](#)
- [Routing Instances for Layer 2 VPN and Layer 3 VPN on page 470](#)
- [Ethernet Encapsulation for Layer 2 VPN on page 471](#)

---

### MPLS Protocol and Label-Switched Paths

Each PE switch must be configured to support the MPLS protocol. You must also configure label-switched paths (LSPs) at the **[edit protocols mpls]** hierarchy level.

---

### IP Over MPLS for Customer Edge Interfaces

You can configure the customer edge interfaces of the PE switches for IP over MPLS using a Layer 3 interface and a static route from the ingress PE switch to the egress PE switch. See “[Configuring MPLS on Provider Edge Switches](#)” on page 188.

---

### BGP Layer 3 VPN Configuration

If you are implementing a Layer 3 virtual private network (VPN), you must configure the BGP routing protocol on the PE switches.

---

### Routing Instances for Layer 3 VPN

If you are implementing a Layer 3 VPN, you must configure a routing instance. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables.

QFX Series and EX4600 devices support VPN routing and forwarding (VRF) routing instances for Layer 3 VPNs.

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table will be **my-instance.inet.0**. All routes for **my-instance** are installed in **my-instance.inet.0**.

---

### Routing Instances for Layer 2 VPN and Layer 3 VPN

If you are implementing a Layer 2 VPN or a Layer 3 VPN, you must configure a routing instance. A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables.

QFX Series devices support the following types of routing instances:

- Layer 2 VPN—To support a Layer 2 VPN
- VPN routing and forwarding (VRF)—To support a Layer 3 VPN

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table will be **my-instance.inet.0**. All routes for **my-instance** are installed in **my-instance.inet.0**.

### Ethernet Encapsulation for Layer 2 VPN

If you are implementing a Layer 2 VPN, you must also configure the physical layer encapsulation type on the customer edge interface and within the routing instance.

## Provider Switch

You must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets. You can add provider switches without changing the configuration of the PE switches.

A provider switch does not analyze packets. It refers to an MPLS label forwarding table and swaps one label for another. The new label determines the next hop along the MPLS tunnel. A provider switch cannot perform push or pop operations.

## Components Required for All Switches in the MPLS Network

The following MPLS components are configured on both the PE switches and the provider switches:

- [Interior Gateway Protocol on page 471](#)
- [Traffic Engineering on page 472](#)
- [MPLS Protocol on page 472](#)
- [RSVP on page 472](#)
- [Family mpls on page 472](#)

### Interior Gateway Protocol

MPLS works in coordination with OSPF as the interior gateway protocol (IGP). Therefore, you must configure OSPF as the IGP on the loopback interface and CE-facing interfaces of both the PE switches and the provider switches.

The CE-facing interfaces can be either Gigabit Ethernet or 10-Gigabit Ethernet interfaces, and they can be configured as either individual interfaces or as aggregated Ethernet interfaces.



**NOTE:** The CE-facing interfaces cannot be configured with VLAN tagging or a VLAN ID. When you configure them to belong to family **mpls**, they are removed from the default VLAN if they were members of that VLAN. They operate as an exclusive tunnel for MPLS traffic.

## Traffic Engineering

---

Traffic engineering maps traffic flows onto an existing physical topology and provides the ability to move traffic flow away from the shortest path selected by the IGP and to a potentially less congested physical path across a network.

Traffic engineering enables the selection of specific end-to-end paths to send given types of traffic through your network. You must configure OSPF traffic engineering on the PE switches and the provider switches.

## MPLS Protocol

---

You must enable the MPLS protocol on all switches that participate in the MPLS network and apply it to the core interfaces of both the PE and provider switches. You do not need to apply it to the loopback interface because the MPLS protocol uses the framework established by the RSVP signaling protocol to create LSPs. On the PE switches, the configuration of the MPLS protocol must also include the definition of an LSP.

## RSVP

---

RSVP is a signaling protocol that allocates and distributes labels throughout an MPLS network. RSVP sets up unidirectional paths between the ingress PE switch and the egress PE switch. RSVP makes the LSPs dynamic; it can detect topology changes and outages and establish new LSPs to allow traffic to move around a failure.

You must enable RSVP and apply it to the loopback interface and the core interface of both the PE and provider switches. The path message contains the configured information about the resources required for the LSP to be established.

When the egress PE switch receives the path message, it sends a reservation message back to the ingress PE switch. This reservation message is passed along from switch to switch along the same path as the original path message. Once the ingress PE switch receives this reservation message, an RSVP path is established.

The established LSP stays active as long as the RSVP session remains active. RSVP continues activity through the transmissions and responses to RSVP path and reservation messages. If the messages stop for three minutes, the RSVP session terminates and the LSP is lost.

RSVP runs as a separate software process in Junos OS and is not in the packet-forwarding path.

## Family mpls

---

You must configure the core interfaces used for MPLS traffic to belong to **family mpls**.



**NOTE:** You can enable **family mpls** on either individual interfaces or on aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

---



**Related Documentation**

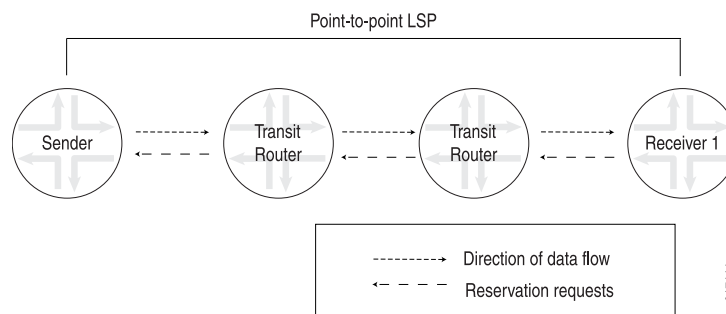
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)
- [Understanding Using MPLS-Based Layer 3 VPNs on Switches on page 162](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 158](#)
- [Configuring MPLS on Provider Edge Switches on page 188](#)
- [Configuring MPLS on Provider Switches on page 192](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers on page 199](#)
- [Configuring a Global MPLS EXP Classifier on page 184](#)
- [Configuring Ethernet over MPLS \(L2 Circuit\) on page 180](#)
- *Junos OS MPLS Applications Library for Routing Devices*
- *Junos OS VPNs Library for Routing Devices*

## RSVP Overview

The RSVP protocol is used by routers to deliver quality-of-service (QoS) requests to all nodes along data flow path(s) and to establish and maintain state for the requested service. RSVP requests generally result in resource reservations in each node along the data path. RSVP has the following attributes:

- Makes resource reservations for unidirectional data flows.
- Allows the receiver of a data flow to initiate and maintain the resource reservation used for that flow, as shown in [Figure 14 on page 473](#).
- Maintains a soft state in routers and hosts, providing graceful support for dynamic membership changes and automatic adaptation to routing changes.
- Depends upon present and future routing protocols, but is not a routing protocol itself.
- Provides several reservation models or styles to fit a variety of applications.
- Supports both IPv4 and IPv6. Note, you can configure the Junos OS to tunnel IPv6 over an MPLS-based IPv4 network. For more information, see the *MPLS Applications Feature Guide for Routing Devices*.

**Figure 14: RSVP Reservation Request and Data Flow**



## MTU Signaling in RSVP

---

The maximum transmission unit (MTU) is the largest size packet or frame, in bytes, that can be sent in a network. An MTU that is too large might cause retransmissions. Too small an MTU might cause the router to send and handle relatively more header overhead and acknowledgments. There are default values for MTUs associated with various protocols. You can also explicitly configure an MTU on an interface.

When an LSP is created across a set of links with different MTU sizes, the ingress router does not know what the smallest MTU is on the LSP path. By default, the MTU for an LSP is 1,500 bytes.

If this MTU is larger than the MTU of one of the intermediate links, traffic might be dropped, because MPLS packets cannot be fragmented. Also, the ingress router is not aware of this type of traffic loss, because the control plane for the LSP would still function normally.

To prevent this type of packet loss in MPLS LSPs, you can configure MTU signaling in RSVP. This feature is described in RFC 3209. Juniper Networks supports the Integrated Services object for MTU signaling in RSVP. The Integrated Services object is described in RFCs 2210 and 2215. MTU signaling in RSVP is disabled by default.

To avoid packet loss due to MTU mismatches, the ingress router needs to do the following:

- Signal the MTU on the RSVP LSP—To prevent packet loss from an MTU mismatch, the ingress router needs to know what the smallest MTU value is along the path taken by the LSP. Once this MTU value is obtained, the ingress router can assign it to the LSP.
- Fragment packets—Using the assigned MTU value, packets that exceed the size of the MTU can be fragmented into smaller packets on the ingress router before they are encapsulated in MPLS and sent over the RSVP-signaled LSP.

Once both MTU signaling and packet fragmentation have been enabled on an ingress router, any route resolving to an RSVP LSP on this router uses the signaled MTU value. For information about how to configure this feature, see *Configuring MTU Signaling in RSVP*.

The following sections describe how MTU signaling in RSVP works:

- *How the Correct MTU Is Signaled in RSVP*
- *Determining an Outgoing MTU Value*
- *MTU Signaling in RSVP Limitations*

## Tunneling LDP LSPs in RSVP LSPs

---

You can tunnel LDP LSPs over RSVP LSPs. The following sections describe how tunneling of LDP LSPs in RSVP LSPs works:

- [Tunneling LDP LSPs in RSVP LSPs Overview on page 5](#)
- [Label Operations on page 6](#)

## Tunneling LDP LSPs in RSVP LSPs Overview

If you are using RSVP for traffic engineering, you can run LDP simultaneously to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. LDP effectively treats the traffic-engineered LSPs as single hops.

When you configure the router to run LDP across RSVP-established LSPs, LDP automatically establishes sessions with the router at the other end of the LSP. LDP control packets are routed hop-by-hop, rather than carried through the LSP. This routing allows you to use simplex (one-way) traffic-engineered LSPs. Traffic in the opposite direction flows through LDP-established LSPs that follow unicast routing rather than through traffic-engineered tunnels.

If you configure LDP over RSVP LSPs, you can still configure multiple OSPF areas and IS-IS levels in the traffic engineered core and in the surrounding LDP cloud.



**NOTE:** Beginning with Junos OS Release 15.1, multi-instance support is extended to LDP over RSVP tunneling for a virtual router routing instance. This allows splitting of a single routing and MPLS domain into multiple domains so that each domain can be scaled independently. BGP labeled unicast can be used to stitch these domains for service FECs. Each domain uses intra-domain LDP over RSVP LSP for MPLS forwarding.

### Related Documentation

- [Label Operations on page 6](#)
- [Configuring a Hierarchy of RSVP LSPs to Tunnel Multiple RSVP LSPs Over a Single RSVP LSP](#)

## Configuring MPLS on Provider Edge Switches

To implement MPLS, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch—and at least one provider switch. You can configure the customer edge (CE) interfaces on the PE switches of the MPLS network using IP over MPLS.

This topic describes how to configure an ingress PE switch and an egress PE switch using IP over MPLS:

1. [Configuring the Ingress PE Switch on page 476](#)
2. [Configuring the Egress PE Switch on page 477](#)

## Configuring the Ingress PE Switch

To configure the ingress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.10.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.6.1/24
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:



**NOTE:** You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
```

3. Configure OSPF traffic engineering:

```
[edit protocols ospf]
user@switch# set traffic-engineering
```

4. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols rsvp]
user@switch# set interface lo0.0
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

5. Configure MPLS traffic engineering.

```
[edit protocols mpls]
user@switch# set traffic-engineering
```

6. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```

7. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
```

8. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 121.100.10.1/16
```

9. Configure this Layer 3 customer edge interface for the routing protocol:

- ```
[edit]
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3.0
```
10. Configure an LSP on the ingress PE switch (192.168.10.1) to send IP packets over MPLS to the egress PE switch (192.168.12.1):
 

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_1 to 192.168.12.1
```
  11. Disable constrained-path LSP computation for this LSP:
 

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_1 no-cspf
```
  12. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:
 

```
[edit routing-options]
user@switch# set static route 2.2.2.0/24 next-hop 192.168.10.1
user@switch# set static route 2.2.2.0/24 resolve
```

## Configuring the Egress PE Switch

To configure the egress PE switch:

1. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 192.168.12.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.20.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.21.1/24
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure OSPF on the loopback interface and the core interfaces:



**NOTE:** You can use the switch address as an alternative to the loopback interface.

- ```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
```
3. Configure RSVP on the loopback interface and the core interfaces:
 

```
[edit protocols rsvp]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface xe-0/0/5.0
user@switch# set rsvp interface xe-0/0/6.0
```
  4. Configure MPLS on the core interfaces:
 

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
```
  5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
```

6. Configure a customer edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit interfaces]
user@switch# set xe-0/0/3 unit 0 family inet address 2.2.2.1/16
```

7. Configure this Layer 3 customer edge interface for the routing protocol:

```
[edit]
user@switch# set protocols ospf area 0.0.0 interface xe-0/0/3
```

8. Configure an LSP on the egress PE switch (192.168.12.1) to send IP packets over MPLS to the ingress PE switch (192.168.10.1):

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_2 to 192.168.10.1
```

9. Disable constrained-path LSP computation for this LSP:

```
[edit protocols mpls]
user@switch# set label-switched-path lsp_2 no-cspf
```

10. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:

```
[edit routing-options]
user@switch# set static route 121.121.121.0/24 next-hop 192.168.12.1
user@switch# set static route 121.121.121.0/24 resolve
```

**Related  
Documentation**

- [MPLS Configuration Guidelines on page 220](#)
- [Configuring MPLS on Provider Switches on page 192](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 146](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 158](#)

## Configuring MPLS on Provider Switches

To implement MPLS, you must configure at least one provider switch as a transit switch for the MPLS packets.

MPLS requires the configuration of an interior gateway protocol (OSPF) and a signaling protocol (RSVP) on the core interfaces and the loopback interface of all the switches. This procedure includes the configuration of OSPF on the provider switch.

To configure the provider switch, complete the following tasks:

1. Configure OSPF on the loopback and core interfaces:



**NOTE:** You can use the switch address as an alternative to the loopback interface.

```
[edit protocols ospf]
user@switch# set area 0.0.0.0 interface lo0.0
user@switch# set area 0.0.0.0 interface xe-0/0/5.0
user@switch# set area 0.0.0.0 interface xe-0/0/6.0
user@switch# set area 0.0.0.0 interface ae0
```



**NOTE:** You cannot use routed VLAN interfaces (RVIs) or Layer 3 subinterfaces as core interfaces.

2. Configure MPLS on the core interfaces:

```
[edit protocols mpls]
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
user@switch# set interface ae0
```

3. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols rsvp]
user@switch# set interface lo0.0
user@switch# set interface xe-0/0/5.0
user@switch# set interface xe-0/0/6.0
user@switch# set interface ae0
```

4. Configure an IP address for the loopback interface and the core interfaces:

```
[edit interfaces]
user@switch# set lo0 unit 0 family inet address 127.1.1.1/32
user@switch# set xe-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set xe-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switch# set ae0 unit 0 family inet address 10.1.9.2/24
```

5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit interfaces]
user@switch# set xe-0/0/5 unit 0 family mpls
user@switch# set xe-0/0/6 unit 0 family mpls
user@switch# set ae0 unit 0 family mpls
```



**NOTE:** You can configure **family mpls** on either individual interfaces or aggregated Ethernet interfaces. You cannot configure it on tagged VLAN interfaces.

#### Related Documentation

- [Configuring MPLS on Provider Edge Switches on page 188](#)
- [MPLS Configuration Guidelines on page 220](#)
- [MPLS Feature Support on QFX Series and EX4600 Switches on page 134](#)
- [Understanding MPLS Components for QFX Series and EX4600 Switches on page 146](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules on page 158](#)

## Verifying That MPLS Is Working Correctly

To verify that MPLS is working correctly, perform the following tasks:

1. [Verifying the Physical Layer on the Switches on page 480](#)
2. [Verifying the Routing Protocol on page 481](#)
3. [Verifying the Core Interfaces Being Used for the MPLS Traffic on page 481](#)
4. [Verifying RSVP on page 481](#)

### Verifying the Physical Layer on the Switches

**Purpose** Verify that the interfaces are up. Perform this verification task on each of the switches.

**Action** user@switch> **show interfaces xe-\* terse**

Interface	Admin	Link	Proto	Local	Remote
xe-0/0/0	up	up			
xe-0/0/0.0	up	up			
xe-0/0/1.0	up	up			
xe-0/0/2.0	up	up			
xe-0/0/3.0	up	up	inet	2.2.2.1/16	
xe-0/0/4.0	up	up			
xe-0/0/5.0	up	up	inet mpls	10.1.5.1/24	
xe-0/0/6.0	up	up	inet mpls	10.1.6.1/24	

**Meaning** The **show interfaces terse** command displays status information about the 10-Gigabit Ethernet interfaces on the switch. This output verifies that the interfaces are **up**. The output for the protocol family (Proto column) of the core interfaces (xe-0/0/5.0 and xe-0/0/6.0), shows that these interfaces are configured as both **inet** and **mpls**. The **Local** column for the core interfaces shows the IP address configured for these interfaces.



## Verifying the Routing Protocol

**Purpose** Verify the state of the configured routing protocol. You should perform this verification task on each of the switches. The state should be **Full**. If you have configured OSPF as the routing protocol, use the **show ospf neighbor** command to verify that the routing protocol is communicating with the switch neighbors.

**Action** user@switch> **show ospf neighbor**

Address	Interface	State	ID	Pri	Dead
127.1.1.1	xe-0/0/5	Full	10.10.10.10	128	39

**Meaning** The **show ospf neighbor** command displays the status of the routing protocol that has been configured on this switch. The output shows that the state is **Full**, meaning that the routing protocol is operating correctly—that is, hello packets are being exchanged between directly connected neighbors. For additional information on checking and monitoring routing protocols, see the [Junos OS Routing Protocols and Policies Command Reference](#).

## Verifying the Core Interfaces Being Used for the MPLS Traffic

**Purpose** Verify that the state of the MPLS interface is **Up**. You should perform this verification task on each of the switches.

**Action** user@switch> **show mpls interface**

Interface	State	Administrative groups
ge-0/0/5	Up	<none>
ge-0/0/6	Up	<none>

**Meaning** The **show mpls interface** command displays the status of the core interfaces that have been configured to belong to **family mpls**. This output shows that the interface configured to belong to **family mpls** is up.

## Verifying RSVP

**Purpose** Verify the state of the RSVP session. You should perform this verification task on each of the switches.

```
user@switch> show mpls session
```

```
Ingress RSVP: 1 sessions
To          From          State  Rt  Style Labelin Labelout LSPname
127.1.1.3   127.1.1.1   Up     0   1 FF      -    300064 lsp_to_pe2_ge1
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
To          From          State  Rt  Style Labelin Labelout LSPname
127.1.1.1   127.1.1.3   Up     0   1 FF  299968    -    lsp_to_pe1_ge1
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

**Meaning** This output confirms that the RSVP sessions are up.

**Related Documentation**

- [Configuring MPLS on Provider Edge Switches on page 188](#)
- [Configuring MPLS on Provider Switches on page 192](#)

## Dynamic Bandwidth Management Using Container LSP Overview

RSVP LSPs with the autobandwidth feature are increasingly deployed in networks to meet traffic engineering needs. However, the current traffic engineering solutions for point-to-point LSPs are inefficient in terms of network bandwidth utilization, mainly because the ingress routers originating the RSVP LSPs either try to fit the LSPs along a particular path without creating parallel LSPs, or do not interact with the other routers in the network and probe for additional available bandwidth.

This feature provides an ingress router with the capability of acquiring as much network bandwidth as possible by creating parallel LSPs dynamically.

- [Understanding RSVP Multipath Extensions on page 482](#)
- [Junos OS RSVP Multipath Implementation on page 483](#)
- [Current Traffic Engineering Challenges on page 484](#)
- [Using Container LSP as a Solution on page 487](#)
- [Junos OS Container LSP Implementation on page 489](#)
- [Configuration Statements Supported for Container LSPs on page 504](#)
- [Impact of Configuring Container LSPs on Network Performance on page 508](#)
- [Supported and Unsupported Features on page 509](#)

### Understanding RSVP Multipath Extensions

The RSVP multipath extensions proposed in the IETF [KOMPELLA-MLSP] allow the setup of traffic engineered multipath label-switched paths (container LSPs). The container LSPs, in addition to conforming to traffic engineering constraints, use multiple independent paths from a source to a destination, thereby facilitating load balancing of

traffic. The multipath extensions require changes to the RSVP-TE protocol and allow for merging of labels at the downstream nodes (similar to LDP), which also helps in preserving forwarding resources.

The multipath extensions to RSVP provide the following benefits:

- Ease of configuration. Typically, multiple RSVP LSPs are configured for either load balancing or bin packing. With a container LSP, there is a single entity to provision, manage, and monitor LSPs. Changes in topology are handled easily and autonomously by the ingress LSP, by adding, changing, or removing member LSPs to rebalance traffic, while maintaining the same traffic engineering constraints.
- RSVP equal-cost multipath (ECMP) inherits the standard benefits of ECMP by absorbing traffic surges.
- Multipath traffic engineering allows for better and complete usage of network resources.
- Knowing the relationship among LSPs helps in computing diverse paths with constraint-based routing. It allows adjustment of member LSPs while other member LSPs continue to carry traffic.
- The intermediate routers have an opportunity to merge the labels of member LSPs. This reduces the number of labels that need to get added to the forwarding plane and in turn reduces the convergence time.

If the number of independent ECMP paths is huge, label merging overcomes the platform limitations on maximum (ECMP) next hops. With point-to-point RSVP LSPs that require link or node protection, the next hops are doubled as each LSP is programmed with both primary and backup next hops. RSVP multipath (or ECMP) obviates the need for backup next hops.

- When there is a link failure, the router upstream to the link failure can distribute traffic from the failed link to the remaining ECMP branches, obviating the need for bypass LSPs. The bypass LSP approach not only requires more state when signaling backup LSPs, but also suffers from scaling issues that result in merge-point timing out a protected path state block (PSB) before point of local repair (PLR) gets a chance to signal the backup LSP.

## Junos OS RSVP Multipath Implementation

In order to deploy RSVP multipath (ECMP) in a network, all the nodes through which ECMP LSPs pass must understand RSVP ECMP protocol extensions. This can be a challenge, especially in a multivendor networks.


Junos OS implements the RSVP multipath extensions without the need for protocol extensions. A single container LSP, which has the characteristics of ECMP and RSVP TE, is provisioned. A container LSP consists of several member LSPs and is set up between the ingress and egress routing device. Each member LSP takes a different path to the same destination. The ingress routing device is configured with all the required parameters to compute the RSVP ECMP LSP. The parameters configured to compute a set of RSVP point-to-point LSPs can be used by the ingress routing device to compute the container LSP as well.

## Current Traffic Engineering Challenges

The main challenge for traffic engineering is to cope with the dynamics of both topology and traffic demands. Mechanisms are needed that can handle traffic load dynamics in scenarios with sudden changes in traffic demand and dynamically distribute traffic to benefit from available resources.

Figure 15 on page 484 illustrates a sample network topology with all the LSPs having the same hold and setup priorities, and admission control restricted on the ingress router. All the links are annotated with a tuple (cost and capacity).

Figure 15: Sample Topology



g042458

Some of the traffic engineering problems seen in [Figure 15 on page 484](#) are listed here:

- **Bin Packing**

This problem arises because of a particular order in which LSPs are signaled. The ingress routers might not be able to signal some LSPs with required demands although bandwidth is available in the network, leading to under-utilization of link capacity.

For example, the following LSPs arrive in the sequence mentioned in [Table 58 on page 485](#).

**Table 58: LSP Sequence Order for Bin Packing**

Time	Source	Destination	Demand	ERO
1	A	E	5	A-C-D-E
2	B	E	10	No ERO

The LSP originating at Router B is not routable as constraint-based routing fails to find a feasible path. However, if Router B is signaled first, both the LSPs are routable. Bin packing happens because of lack of visibility of individual per-LSP, per-device bandwidth demands at the ingress routing device.

Bin packing can also happen when there is no requirement for ordering of LSPs. For example, if there is an LSP with demand X and there are two different paths to the destination from the ingress router with available bandwidths Y1 and Y2, such that Y1 is less than X, Y2 is less than X, and Y1 plus Y2 is greater than or equal to X.

In this case, even though there are enough network resources in terms of available bandwidth to satisfy the aggregate LSP demand X, the LSP might not be signaled or re-optimized with the new demand. In [Figure 15 on page 484](#), with container LSP support, the ingress B creates two LSPs each of size 5 when demand 10 is posed. One LSP is routed along B-C-E and another one along B-C-D-E.

- **Deadlock**

Considering [Figure 15 on page 484](#), the LSPs follow the sequence mentioned in [Table 59 on page 485](#).

**Table 59: LSP Sequence Order for Deadlock**

Time	Source	Destination	Demand	ERO	Event
1	A	E	2	A-C-D-E	Constraint-based routing with RSVP signaling
2	B	E	2	B-C-D-E	Constraint-based routing with RSVP signaling
3	A	E	2 to 20	A-C-D-E	Constraint-based routing fails, no RSVP signaling

At time 3, the demand on LSP from A to E increases from 2 to 20. If autobandwidth is configured, the change does not get detected until the adjustment timer expires. In

the absence of admission control at A, the increased traffic demand might cause traffic to drop on other LSPs that share common links with the mis-behaving LSP.

This happens due to the following reasons:

- Lack of global state at all the ingress routers
- Signaling of mis-behaving demands
- Tearing down of mis-behaving demands

With container LSP configured, ingress A has more chances of splitting the load (even incrementally if not fully) across multiple LSPs. So, LSP from A is less likely to see prolonged traffic loss.

- **Latency Inflation**

Latency inflation is caused by the autobandwidth and other LSPs parameters. Some of the other factors that contribute to latency inflation include:

- LSP priority

LSPs choose longer paths because shorter paths between data centers located in the same city can be congested. The bandwidth on the shorter paths can get exhausted by equal or higher priority LSPs. Due to periodic LSP optimization by autobandwidth, LSP can get rerouted to a higher delay path. When many LSPs undergo less than optimal path selection, they can potentially form a chain of dependencies. Modifying the LSP priorities dynamically is a workaround to the issue; however, dynamically adjusting LSP priorities to find shorter paths is a challenging task.

- All or Nothing policy

When the demand on an LSP increases and at least one of the links along the shorter path is close to its reservation limit, LSP optimization can force the LSP to move to a longer latency path. LSP has to traverse a long path even though the short path is capable of carrying most of the traffic.

- Minimum and maximum bandwidth

Minimum and maximum bandwidth specify the boundaries for LSP sizes. If minimum bandwidth is small, an LSP is more prone to autobandwidth adjustment because a small change in bandwidth is enough to cross the threshold limits. LSPs might reroute although bandwidth is available. On the other hand, if the minimum bandwidth is large, network bandwidth might be wasted. If the maximum bandwidth value is small, a large number of LSPs might be needed at the ingress router to accommodate the application demand. If the maximum bandwidth is large, the LSPs can grow larger in size. Such LSPs can suffer because of an all or nothing policy.

- Autobandwidth adjustment threshold

Bandwidth threshold dictates if LSPs need to be re-optimized and resized. If the value is small, LSPs are frequently re-optimized and rerouted. That might cause CPU spike because applications or protocols, such as BGP resolving over the LSPs, might keep the Routing Engine busy doing next-hop resolution. A large value might make an LSP immobile. With container LSP configured, an LSP is less likely to get subjected

to one or no policy. An ingress router originates multiple LSPs, although not all LSPs potentially traverse high latency paths.

- **Predictability**

Service providers often want predictable behavior in terms of how LSPs get signaled and routed. Currently, without any global coordination, it is difficult to set up the same set of LSPs in a predictable way. Consider the two different orderings in [Table 60 on page 487](#) and [Table 61 on page 487](#). The ERO that an LSP uses depends on its signaling time.

**Table 60: LSP Sequence Order for Predictability**

Time	Source	Destination	Demand	ERO
1	A	E	5	A-C-D-E
2	B	E	5	B-C-E

**Table 61: LSP Sequence Order for Predictability**

Time	Source	Destination	Demand	ERO
1	B	E	5	B-C-E
2	A	E	5	A-C-D-E

Container LSP does not directly help LSPs find predictable EROs. If LSPs are getting rerouted because of an all or no policy without container LSP configured, such LSPs might see less churn if container LSPs are configured, because smaller LSPs have better chances of finding a shorter or same path.

## Using Container LSP as a Solution

A container LSP can be used as a solution to the challenges faced by the current traffic engineering features. Considering [Figure 15 on page 484](#), when the demand X on a container LSP increases with the network capacity (max-flow) being more than the demand, the following approaches come into effect with a container LSP:

- [Accommodating the New Demand X on page 487](#)
- [Creating New LSPs to Meet Demand X on page 488](#)
- [Assigning Bandwidth to the New LSPs on page 488](#)
- [Controlling the LSP Paths on page 488](#)

### Accommodating the New Demand X

In the current implementation, autobandwidth attempts to re-signal an LSP with the new demand X and follows the all or nothing policy as mentioned earlier.

The container LSP approach computes several small (smaller than demand X) bandwidth LSPs such that the aggregate bandwidth is not less than X, and the ingress router performs

this adjustment periodically. One of the triggers to create new LSPs or to delete old LSPs can be changed in aggregate bandwidth. The ingress router then load-balances the incoming traffic across the newly created LSPs.

### Creating New LSPs to Meet Demand X

---

Although the number of new LSPs created can be a maximum of the allowed configurable limit, there is not much benefit from these LSPs once the number of LSPs exceeds the number of possible diverse paths or equal-cost multipaths (ECMPs). The benefit of creating the smaller LSPs is seen when an ingress router uses the newly created LSPs for load-balancing traffic. This, however, depends on the network topology and state.

Creating multiple parallel LSPs by all the ingress routers in the network can lead to scaling issues at the transit routers. Thus, the number of new LSPs to be created depends on the size of the individual LSPs and the given aggregate demand, X in this case.

### Assigning Bandwidth to the New LSPs

---

In general, there can be a number of heuristics to allocate bandwidths to the newly created LSPs. An ingress router can solve an optimization problem in which it can maximize a given utility function. The output of an optimization problem is assigning optimal bandwidth values. However, to solve an optimization problem, the number of newly created LSPs has to be fixed. Therefore, it is complex to optimize the number and size of each LSP. Thus, to simplify the problem, the same amount of bandwidth is assumed for all the newly created LSPs, and then the number of required LSPs is computed.

### Controlling the LSP Paths

---

The flexibility to control the LSP paths is expressed in terms of the configuration for point-to-point LSPs and container LSPs. Controlling the LSP paths using the configuration parameters can be applied under two different aspects:

- **Topology**—There are no topology constraints with this feature. Each member LSP is treated like a point-to-point LSP and is re-optimized individually. An ingress router does not try to compute equal IGP cost paths for all its LSPs, but instead it computes paths for all the LSPs using current traffic engineering database information. While computing a path, constraint-based routing adheres to any constraints specified through the configuration, although there is no change in the constraint-based routing method for path computation.
- **When to create a new LSP**—When to create a new LSP can be explicitly specified. By default, an ingress router periodically computes the aggregate traffic rate by adding up the traffic rate of all the individual LSPs. Looking at the aggregate bandwidth and configuration, the ingress router recomputes the number of LSPs and the bandwidths of the LSPs. The new LSPs are then signaled or the existing LSPs are re-signaled with the updated bandwidth. Instead of looking at the instantaneous aggregate rate, the ingress routers can compute an average (of aggregates) over some duration by removing outlier samples (of aggregates). Managing the LSPs that remain outstanding and active by considering aggregate bandwidth is more scalable than creating the new LSPs based on the usage of a particular LSP. The intervals and thresholds can be configured to track the aggregate traffic and trigger adjustment. These dynamic LSPs co-exist and interoperate with per-LSP autobandwidth configuration.



## Junos OS Container LSP Implementation

A container LSP is an ECMP TE LSP that acts like a container LSP consisting of one or more member LSPs. A point-to-point TE LSP is equivalent to a container LSP with a single member LSP. Member LSPs are added to the container LSP through a process called splitting, and removed from the container LSP through a process called merging.

- [Container LSP Terminology on page 489](#)
- [LSP Splitting on page 490](#)
- [LSP Merging on page 492](#)
- [Node and Link Protection on page 494](#)
- [Naming Convention on page 494](#)
- [Normalization on page 495](#)
- [Constraint-Based Routing Path Computation on page 500](#)
- [Sampling on page 501](#)
- [Support for NSR, IPG-FA, and Static Routes on page 501](#)

### Container LSP Terminology

The following terms are defined in the context of a container LSP:

- **Normalization**—An event occurring periodically when an action is taken to adjust the member LSPs, either to adjust their bandwidths, their number, or both. A normalization process is associated with a sampling process and periodically estimates aggregate utilization of a container LSP.
- **Nominal LSP**—The instance of a container LSP that is always present.
- **Supplementary LSP**—The instances or sub-LSPs of a container LSP, which are dynamically created or removed.

Autobandwidth is run over each of the member LSPs, and each LSP is resized according to the traffic it carries and the autobandwidth configuration parameters. The aggregate demand on a container LSP is tracked by adding up the bandwidth across all the member LSPs.

- **Minimum signaling-bandwidth**—The minimum bandwidth with which a member LSP is signaled at the time of normalization or initialization. This could be different from the minimum-bandwidth defined under autobandwidth.
- **Maximum signaling-bandwidth**—The maximum bandwidth with which a member LSP is signaled at the time of normalization or initialization. This could be different from the maximum-bandwidth defined under autobandwidth.
- **Merging-bandwidth**—Specifies the lower bandwidth threshold on the aggregate bandwidth usage, such that if the aggregate usage falls below this value, the ingress router merges the member LSPs at the time of normalization.
- **Splitting-bandwidth**—Specifies the upper bandwidth threshold on the aggregate bandwidth usage, such that if the aggregate usage exceeds this value, the ingress router splits the member LSPs at the time of normalization.

- **Aggregate minimum-bandwidth**—Sum of merging-bandwidth of the current active member LSPs. This minimum bandwidth is different from the autobandwidth minimum-bandwidth.
- **Aggregate maximum-bandwidth**—Sum of the splitting-bandwidth of the current active member LSPs. This maximum bandwidth is different from the autobandwidth maximum-bandwidth.

---

## LSP Splitting

- [Operational Overview on page 490](#)
- [Operational Constraints on page 491](#)
- [Supported Criteria on page 491](#)
- [Splitting Triggers on page 492](#)

### *Operational Overview*

The LSP splitting mechanism enables an ingress router to create new member LSPs or to re-signal existing LSPs with different bandwidths within a container LSP when a demand X is placed on the container LSP. With LSP splitting enabled, an ingress router periodically creates a number of LSPs (by signaling new ones or re-signaling existing ones) to accommodate a new aggregate demand X. In the current implementation, an ingress router tries to find an LSP path satisfying a demand X and other constraints. If no path is found, either the LSP is not signaled or it remains up, but with the old reserved bandwidth.

Between two normalization events (splitting or merging), individual LSPs might get re-sigaled with different bandwidths due to the autobandwidth adjustments. If a container LSP is not configured with autobandwidth, then the member LSPs are signaled with the static bandwidth value, if configured. There is no dynamic splitting in this case, as there is no dynamic estimation of aggregate bandwidth. The splitting adjustments with a specific bandwidth value can be manually triggered.



#### NOTE:

Be aware of the following considerations for LSP splitting:

- After LSP splitting, the ingress router continues to inject one forwarding adjacency. Forwarding adjacencies are not supported in IGP for this feature.
  - Between two normalization events, two LSPs might have different bandwidths subjected to autobandwidth constraints.
  - After LSPs are split (or merged), make-before-break uses the fixed filter (FF) style sharing unless the adaptive option is configured. However, two different LSPs do not do the shared explicit (SE) style sharing for this feature.
  - When LSPs are re-sigaled with modified bandwidths, some of the LSPs might not get signaled successfully, leading to failover options.
-

### ***Operational Constraints***

LSP splitting has the following operational constraints:

- LSP bandwidth—Although there are a number of ways to allocate bandwidth values to the LSPs, the Junos OS implementation supports only an equal-bandwidth allocation policy when normalization is done, wherein all the member LSPs are signaled or re-signaled with equal bandwidth.
- Number of LSPs—If an ingress router is configured to have a minimum number of LSPs, it maintains the minimum number of LSPs even if the demand can be satisfied with less than the minimum number of LSPs. In case the ingress router is unable to do constraint-based routing for computations on the sufficient number of LSPs or signal sufficient number of LSPs, the ingress router resorts to a number of fallback options.

By default, an incremental approach is supported as a fallback option (unless configured differently), where an ingress router makes attempts to bring up the sufficient number of LSPs, such that the new aggregate bandwidth exceeds the old aggregate bandwidth (and is as close to the desired demand as possible). The ingress router then load-balances traffic using the LSPs. The LSPs that could not be brought up are removed by the ingress router.

### ***Supported Criteria***

When a container LSP signals a member LSP, the member LSP gets signaled with minimum-signaling-bandwidth. Since each member LSP is configured with autobandwidth, between two normalization events, each LSP can undergo autobandwidth adjustment multiple times. As the traffic demand increases, the ingress router creates additional supplementary LSPs. All member LSPs are used for ECMP, so they should roughly have the same reserved bandwidth after normalization.

For example, if there are K LSPs signaled after normalization, each LSP is signaled with equal bandwidth B. The total aggregate bandwidth reserved is B.K, where B satisfies the following condition:

- Minimum signaling-bandwidth is less than or equal to B, which in turn is less than or equal to the maximum signaling-bandwidth  
(minimum-signaling-bandwidth  $\leq$  B  $\leq$  maximum-signaling-bandwidth)

Until the next normalization event, each member LSP undergoes several autobandwidth adjustments. After any autobandwidth adjustment, if there are N LSPs with reserved bandwidths  $b_i$ , where  $i=1,2,\dots, N$ , each  $b_i$  should satisfy the following condition:

- Minimum bandwidth is less than or equal to  $b_i$ , which in turn is less than or equal to the maximum bandwidth  
(minimum-bandwidth  $\leq b_i \leq$  maximum-bandwidth)

Both the above-mentioned conditions are applicable for per member LSP (nominal and supplementary), and essentially have the reserved bandwidth to exist within a range.

### ***Splitting Triggers***

Every time the normalization timer expires, the ingress router decides if LSP splitting is required. The ingress router works with the aggregate bandwidth instead of the individual LSP bandwidths. The following two variables are defined for aggregate bandwidth:

- **Current-Aggr-Bw**—Sum of reserved bandwidths of all current member LSPs.
- **New-Aggr-Bw**—Sum of traffic rates on all current member LSPs based on sampling.

Taking for example, if there are N member LSPs in the network at the time of normalization, the two approaches to trigger LSP splitting are as follows:

- Absolute trigger—LSP splitting is performed when **New-Aggr-Bw** is greater than **Aggregate-maximum-bandwidth**.

(**New-Aggr-Bw** > **Aggregate-maximum-bandwidth**)

- Relative trigger—The **Current-Aggr-Bw** is compared with **New-Aggr-Bw** at the ingress routing device. LSP splitting is performed when the difference in the bandwidth amount is off by a threshold.

$([1-a] \times \text{Current-Aggr-Bw} < \text{New-Aggr-Bw} < [1+a] \times \text{Current-Aggr-Bw})$ , where  $0 \leq a \leq 1$

When **New-Aggr-Bw** is greater than or equal to  $[1+a]$  multiplied by **Current-Aggr-Bw**, the ingress routing device does not perform normalization, but instead LSP splitting is done. However, when both LSP splitting and LSP merging are configured on the ingress router, LSP splitting is triggered on the ingress router when one of the two conditions is satisfied.

---

### ***LSP Merging***

- [Operational Overview on page 492](#)
- [Operational Constraints on page 493](#)
- [Merging Triggers on page 493](#)

#### ***Operational Overview***

Junos OS supports two kinds of LSPs – CLI-configured LSPs and dynamically created LSPs. The CLI-configured LSPs are created manually and remain in the system until the configuration is modified. The dynamic LSPs are created dynamically by next generation MVPN, BGP virtual private LAN service (VPLS), or LDP, based on a template configuration, and are removed from the system when not used by any application for a certain duration. LSP merging follows a similar approach as dynamic LSPs.

LSP merging enables an ingress routing device to dynamically eliminate some member LSPs of the container LSP so less state information is maintained in the network. If an ingress router provisions several member LSPs between the ingress and egress routers, and there is an overall reduction in aggregate bandwidth (resulting in some LSPs being under-utilized), the ingress router distributes the new traffic load among fewer LSPs.

Although there are a number of ways to merge the member LSPs, Junos OS supports only overall-merge when normalization is being performed. An ingress router considers the aggregate demand and the minimum (or maximum) number of LSPs and revises the number of LSPs that should be active at an ingress routing device. As a result, the following can take place periodically as the normalization timer fires:

- Re-signaling some of the existing LSPs with updated bandwidth
- Creating new LSPs
- Removing some of the existing LSPs

### **Operational Constraints**

If a container LSP is not configured with autobandwidth, then the member LSPs are signaled with the static bandwidth value, if configured. LSP merging does not happen because there is no dynamic estimation of aggregate bandwidth. However, a manual trigger for splitting and adjusting with a specific bandwidth value can be configured.



#### **NOTE:**

- Nominal LSPs are never deleted as part of LSP merging.
- Before deleting an LSP, the LSP is made inactive, so that traffic shifts to other LSPs before removing the LSP. This is because RSVP sends PathTear before deleting routes and next hops from the Packet Forwarding Engine.
- When member LSPs are re-signaled with modified bandwidth, it might happen that some LSPs do not get signaled successfully.

### **Merging Triggers**

Every time the normalization timer expires, the ingress router decides if LSP merging is required. The ingress router works with the aggregate bandwidth instead of the individual LSP bandwidths. The following two variables are defined for aggregate bandwidth:

- **Current-Aggr-Bw**—Sum of reserved bandwidths of all current member LSPs.
- **New-Aggr-Bw**—Sum of traffic rates on all current member LSPs based on sampling.

For example, if there are N member LSPs in the network at the time of normalization, the two approaches to trigger LSP merging are as follows:

- Absolute trigger—LSP merging is performed when **New-Aggr-Bw** is less than **Aggregate-minimum-bandwidth**.  
(**New-Aggr-Bw** < **Aggregate-maximum-bandwidth**)
- Relative trigger—The **Current-Aggr-Bw** is compared with **New-Aggr-Bw** at the ingress routing device. LSP merging is performed when the difference in the bandwidth amount is off by a threshold.  
$$([1-a] \times \text{Current-Aggr-Bw} < \text{New-Aggr-Bw} < [1+a] \times \text{Current-Aggr-Bw}, \text{ where } 0 \leq a \leq 1)$$

When the **New-Aggr-Bw** value is less than or equal to  $[1+a]$  multiplied by the **Current-Aggr-Bw** value, the ingress routing device does not perform normalization, but instead LSP merging is done. However, when both LSP splitting and LSP merging are configured on the ingress router, LSP splitting is triggered on the ingress router when one of the two conditions is satisfied.

### Node and Link Protection

---

Junos OS supports the following mechanisms for node and link protection:

- Fast-reroute
- Link protection
- Node-link protection

Only one of the above-mentioned modes of protection can be configured on an ingress routing device at any given time. All member LSPs (nominal and supplementary) use the same mode of protection that is configured.

### Naming Convention

---

While configuring a container LSP, a name is assigned to the LSP. The name of a nominal and a supplementary LSP is formed by adding the configured-name suffix and an auto-generated suffix to the name of the container LSP. The name of the container LSP is unique and is checked for accuracy during the configuration parsing. The container LSP name should uniquely identify parameters, such as the ingress and egress router names.



**NOTE:** A container LSP member LSP and a point-to-point LSP on an ingress routing device cannot have the same LSP name.

The container LSPs follow a number-based LSP naming convention. For example, if the nominal LSP's configured name is **bob** and the number of member LSPs is  $N$ , the member LSPs are named **bob-*<configured-suffix>-1***, **bob-*<configured-suffix>-2***, ..., and **bob-*<configured-suffix>-N***.

After a normalization event, the number of member LSPs can change. For example, if the number of member LSPs increases from six to eight, then the ingress routing device keeps the first six LSPs named **bob-*<configured-suffix>-1***, **bob-*<configured-suffix>-2***, ..., and **bob-*<configured-suffix>-6***. The two additional LSPs are named **bob-7** and **bob-8**. The original LSPs might need to be re-optimized if their signaled bandwidth changes.

Similarly, if the number of member LSPs reduces from eight to six, the ingress routing device re-signals the member LSPs in such a way that the remaining active LSPs in the system are named **bob-*<configured-suffix>-1***, **bob-*<configured-suffix>-2***, ..., and **bob-*<configured-suffix>-6***.

In the process of creating new LSPs, an RSVP LSP named **bob-*<configured-suffix>-7*** can be configured.

## Normalization

- [Operational Overview on page 495](#)
- [Operational Constraints on page 495](#)
- [Inter-Operation with Autobandwidth on page 496](#)

### Operational Overview

Normalization is an event that happens periodically. When it happens, a decision is made on the number of member LSPs that should remain active and their respective bandwidths in a container LSP. More specifically, the decision is made on whether new supplementary LSPs are to be created, or any existing LSPs are required to be re-signaled or deleted during the normalization event.

Between two normalization events, a member LSP can undergo several autobandwidth adjustments. A normalization timer, similar to re-optimization timer, is configured. The normalization timer interval should be no less than the adjustment interval or optimization timer.



**NOTE:** Normalization is not triggered based on network events, such as topology changes.

### Operational Constraints

Normalization has the following operational constraints:

- Normalization happens only when none of the member LSPs are undergoing re-optimization or make-before-break. Normalization starts when all the member LSPs complete their ongoing make-before-break. If normalization is pending, new optimization should not be attempted until the normalization is complete.
- After normalization, an ingress routing device first computes a set of bandwidth-feasible paths using constraint-based routing computations. If enough constraint-based routing computed paths are not brought up with an aggregate bandwidth value that exceeds the desired bandwidth, several failover actions are taken.
- After a set of bandwidth-feasible paths are available, the ingress routing device signals those paths while keeping the original set of paths up with the old bandwidth values. The make-before-break is done with shared explicit (SE) sharing style, and when some of the LSPs do not get successfully re-signaled, a bounded number of retries is attempted for a specified duration. Only when all the LSPs are successfully signaled does the ingress router switch from the old instance of the container LSP to the newer instance. If all LSPs could not be successfully signaled, the ingress router keeps those instances of members that are up with higher bandwidth values.

For example, if the bandwidth of an old instance of a member LSP (LSP-1) is 1G, the LSP is split into LSP-1 with bandwidth 2G and LSP-2 with bandwidth 2G. If the signaling of LSP-1 with bandwidth 2G fails, the ingress router keeps LSP-1 with bandwidth 1G and LSP-2 with bandwidth 2G.

When there is a signaling failure, the ingress routing device stays in the error state, where some LSPs have updated bandwidth values only if the aggregate bandwidth has increased. The ingress router makes an attempt to bring up those LSPs that could not be successfully signaled, resulting in minimum traffic loss.

- If an LSP goes down between two normalization events, it can increase the load on other LSPs that are up. In order to prevent overuse of other LSPs, premature normalization can be configured in case of LSP failure. LSPs can go down because of pre-emption or lack of node or link protection. It might not be necessary to bring up the LSPs that are down because the normalization process re-runs the constraint-based routing path computations.

### ***Inter-Operation with Autobandwidth***

Taking as an example, there is one nominal LSP named LSP-1 configured with the following parameters:

- Splitting-bandwidth and maximum-signaling-bandwidth of 1G
- Merging-bandwidth and minimum-signaling-bandwidth of 0.8G
- Autobandwidth

Normalization is performed differently in the following scenarios:

- [Changes in Per-LSP Autobandwidth Adjustments on page 496](#)
- [Changes in Traffic Growth on page 498](#)
- [Computed Range and Configured Feasible Ranges on page 498](#)

### ***Changes in Per-LSP Autobandwidth Adjustments***

[Table 62 on page 496](#) illustrates how normalization splits and merges member LSPs as autobandwidth adjustments change per-LSP bandwidth with unconditional normalization.

**Table 62: Normalization with Per-LSP Autobandwidth Adjustment Changes**

Normalization Time	Current State	Events	Adjusted State
T0	No state.	Initialization	LSP-1 is signaled with bandwidth of 0.8G
T1	LSP-1 usage increases to 1.5G	<ul style="list-style-type: none"> <li>• Multiple autobandwidth adjustments since T0 is possible.</li> <li>• The ingress router decides to split LSP-1 into two LSPs, and creates LSP-2.</li> </ul>	LSP-1 = 0.8G LSP-2 = 0.8G
T2	LSP-1 usage increase to 2G LSP-2 usage increases to 0.9G (within limits)	<ul style="list-style-type: none"> <li>• Aggregate bandwidth is 2.9G, which exceeds aggregate splitting maximum of 2G.</li> <li>• The ingress router decides to split LSP-1 into three LSPs, and creates LSP-3.</li> </ul>	LSP-1 = 1G LSP-2 = 1G LSP-3 = 1G



Table 62: Normalization with Per-LSP Autobandwidth Adjustment Changes (*continued*)

Normalization Time	Current State	Events	Adjusted State
T3	LSP-3 usage increases to 1.5G	<ul style="list-style-type: none"> <li>Aggregate bandwidth is 3.5G with a maximum aggregate splitting of 3G.</li> <li>The ingress router decides to split LSP-1 into four LSPs, and creates LSP-4.</li> </ul>	LSP-1 = 1G LSP-2 = 1G LSP-3 = 1G LSP-4 = 1G
T4	LSP-2 usage drops to 0.5G	<ul style="list-style-type: none"> <li>Aggregate bandwidth is 3G.</li> <li>The ingress router decides to merge LSP-1 and removes LSP-4.</li> </ul>	LSP-1 = 1G LSP-2 = 1G LSP-3 = 1G

Because autobandwidth is configured on a per-LSP basis, every time there is an autobandwidth adjustment, the ingress router re-signals each LSP with **Max Avg Bw**.

Another approach to handling the changes in per-LSP autobandwidth adjustments is to not allow individual LSPs to run autobandwidth on the ingress router, but to run autobandwidth in passive (monitor) mode. This way, sampling is done at every statistics interval for member LSPs only, and normalization is performed for the container LSP alone instead of acting on individual LSPs adjustment timer expiry.

As a result, the number of re-signaling attempts and bandwidth fluctuations for a given member LSP is reduced. Only the computed bandwidth-values per-member LSP is used by the ingress router to find an aggregate bandwidth to be used during normalization. Configuring autobandwidth adjustment followed by normalization (adjustments and normalization intervals are comparable) can lead to considerable overhead because of re-signaling.

Taking the same example, and applying the second approach, LSP-1 goes from 0.8G to 1.5G and then back to 0.8G. If the normalization timer is of the same order as the adjustment interval, the ingress router leaves LSP-1 alone with its original 0.8G and only signals LSP-2 with 0.8G. This helps achieve the final result of normalization, thus avoiding the extra signaling attempt on LSP-1 with 1.5G at adjustment timer expiry.

Because member LSPs always use equal bandwidth, any adjustment done on member LSPs is undone. The member LSPs are re-signaled with reduced bandwidth when compared to the reserved capacity in adjustment trigger with normalization trigger. Therefore, avoiding adjustment trigger for member LSPs might be useful assuming that normalization and adjustment intervals are of the same order.



**NOTE:** We recommend that the normalization timer be higher than the autobandwidth adjustment interval and regular optimization duration, as the traffic trends are observed at a longer time scale and normalization is performed one-to-three times per day. An LSP can undergo optimization for the following reasons:

- Normal optimization
- Autobandwidth adjustment
- Normalization

### *Changes in Traffic Growth*

Table 63 on page 498 illustrates how normalization is performed when traffic grows in large factor.

**Table 63: Normalization with Traffic Growth**

Normalization Time	Current State	Events	Adjusted State
T0	No state		LSP-1 is signaled with bandwidth of 0.8G
T1	LSP-1 usage increase to 3G	<ul style="list-style-type: none"> <li>• Aggregate usage exceeds maximum splitting bandwidth</li> <li>• The ingress router decides to split LSP-1, and creates two more supplementary LSPs</li> </ul>	LSP-1 = 1G LSP-2 = 1G LSP-3 = 1G

Having fewer LSPs is preferred over signaling four LSPs each with 0.8G bandwidth, unless there is a constraint on the minimum number of LSPs.

### *Computed Range and Configured Feasible Ranges*

When an ingress router is configured with the minimum and maximum number of LSPs, and per LSP splitting-bandwidth and merging-bandwidth values, the bandwidth thresholds are used for splitting and merging. For this, the number of LSPs (N) should satisfy the following constraints:

$$\text{minimum-member-lsps} \leq N \leq \text{maximum-member-lsps}$$

At the time of normalization, based on the aggregate demand X:

$$\lceil X/\text{splitting-bandwidth} \rceil \leq N \leq \lfloor X/\text{merging-bandwidth} \rfloor$$

The above-mentioned constraints provide two ranges for N to work from. If the two ranges for N are overlapping, N will be selected from the overlapping interval (lowest possible N) to keep the number of LSPs small in the network.

Otherwise, if maximum-member-lsps is less than  $\lceil X/\text{splitting-bandwidth} \rceil$ , the ingress router keeps (at maximum) the maximum-member-lsps in the system, and the bandwidth

of each LSP is  $\lceil X/\text{maximum-member-lsps} \rceil$  or the maximum-signaling-bandwidth, whichever is less. It is possible that some LSPs might not get signaled successfully.

Similarly, if minimum-member-lsps is greater than  $\lceil X/\text{merging-bandwidth} \rceil$ , the ingress router keeps (at minimum) the minimum-member-lsps in the system, and the bandwidth of each LSP is  $\lceil X/\text{minimum-member-lsps} \rceil$  or the minimum-signaling-bandwidth, whichever is less.

Taking as an example, normalization is performed as following in these cases:

- Case 1
  - minimum-member-lsps = 2
  - maximum-member-lsps = 10
  - aggregate demand = 10G
  - merging-bandwidth = 1G
  - splitting-bandwidth = 2.5G

In this case, the ingress routing device signals four member LSPs each with a bandwidth of 2G.

- Case 2
  - minimum-member-lsps = 5
  - maximum-member-lsps = 10
  - aggregate demand = 10G
  - merging-bandwidth = 2.5G
  - splitting-bandwidth = 10G

In this case, the ingress routing device signals five member LSPs each with a bandwidth of 2G. Here, the static configuration on the number of member LSPs takes precedence.

- Case 3
  - minimum-signaling-bandwidth = 5G
  - maximum-signaling-bandwidth = 40G
  - merging-bandwidth = 10G
  - splitting-bandwidth = 50G

When a container LSP comes up, the nominal LSP is signaled with minimum-signaling-bandwidth. At the time of normalization, the new-aggregate-bandwidth is 100G. To find N and the bandwidth of each LSP, N should satisfy the following constraint:

$$100/50 \leq N \leq 100/10, \text{ which gives } 2 \leq N \leq 10$$

Therefore, N is equal to:

- $N = 2$ , bandwidth =  $\min \{100/2G, 40G\} = 40G$

This option does not satisfy the new aggregate of 100G.

- $N = 3$ , bandwidth =  $\min \{100/3G, 40G\} = 33.3G$

This option makes the aggregate bandwidth equal to 100G.

In this case, the ingress routing device signals three LSPs each with a bandwidth of 33.3G.



**NOTE:** The ingress router does not signal an LSP smaller than the minimum-signaling-bandwidth.

---

### Constraint-Based Routing Path Computation

Although there are no changes in the general constraint-based routing path computation, with a container LSP, there is a separate module that oversees the normalization process, schedules constraint-based routing events, and schedules switchover from an old instance to a new instance, when appropriate. An ingress routing device has to handle the constraint-based routing path computation periodically. When normalization occurs, an ingress router has to compute constraint-based routing paths, if the number of LSPs or the bandwidth of the LSPs needs to be changed.

For example, there are K LSPs at the ingress router with bandwidth values X-1, X-2, ..., and X-K. The current aggregate bandwidth value is Y, which is the sum of X-1 plus X-2 plus X-K. If there is a new demand of W, the ingress router first computes how many LSPs are required. If the ingress router only needs N LSPs (LSP-1, LSP-2, ..., and LSP-N) each with bandwidth value B, the task of the constraint-based routing module is to provide a set of bandwidth-feasible LSPs that can accommodate the new aggregate demand which is not less than Y.

The ingress router then tries to see if the constraint-based routing paths can be computed successfully for all N LSPs. If the paths for all the LSPs are found successfully, the constraint-based routing module returns the set to the normalization module.

It is possible that the constraint-based routing computation is not successful for some LSPs. In this case, the ingress routing device takes the following action:

- If the configuration allows for incremental-normalization, implying if the ingress router has enough LSPs whose aggregate exceeds Y, the constraint-based routing module returns that set of paths.
- Whether increment-normalization is configured or not, if constraint-based routing paths could not be computed for a sufficient number of LSPs, the ingress router has to repeat the process of finding a new set of LSPs. Initially, the ingress router starts with the lowest value of N from the feasible region. Every time, the ingress router has to revise the number, it linearly increases it by 1. As a result, per LSP bandwidth becomes less and therefore, there is a greater chance of successful signaling. The process is repeated for all feasible values of N (or some bounded number of times or duration as configured).

The ingress router signals the LSPs after successful computations of the constraint-based routing path computation. It might happen that when the LSPs are signaled, signaling of many LSPs fail. In addition to the constraint-based routing path computations to be successful, the RSVP signaling should also succeed, such that the new aggregate is not less than the old aggregate bandwidth.

### Sampling

Sampling is important for normalization to function. With sampling configured, an ingress routing device is able to make a statistical estimate of the aggregate traffic demands. Every time the sampling timer fires, the ingress routing device can consider traffic rates on different LSPs and compute an aggregate bandwidth sample. This sampling timer is different from the statistics sampling done periodically by RSVP on all LSPs. The aggregate bandwidth is a sample to be used at the time of normalization. An ingress routing device can save past samples to compute an average (or some other statistical measure) and use it the next time normalization happens.

To remove any outlier samples, a sampling token is configured. In other words, from all the aggregate samples collected during the configured time, the bottom and top outliers are ignored before computing a statistical measure from the remaining samples.

The following two methods of computing an aggregate bandwidth value are supported:

- **Average**—All the aggregate bandwidth samples are considered by the ingress routing device, and then all the outlier samples are removed. The average bandwidth value is computed from the remaining samples to be used during normalization.
- **Max**—All the aggregate bandwidth samples are considered by the ingress routing device, and then all the outlier samples are removed. The maximum bandwidth value is picked from the remaining samples to be used during normalization.

The time duration, the number of past aggregate samples to store, the percentile value to determine, and the ignore outliers are user-configurable parameters.

### Support for NSR, IPG-FA, and Static Routes

Starting with Junos OS Release 15.1, container label-switched paths (LSPs) provide support for nonstop active routing (NSR), IGP forwarding adjacency (FA), and static routes to address the requirements of wider business cases.

- [NSR Support on page 501](#)
- [IPG-FA Support on page 503](#)
- [Static Route Support on page 504](#)

#### **NSR Support**

A container LSP has the characteristics of ECMP and RSVP traffic engineering. Because a container LSP consists of several member LSPs between an ingress and an egress router, with each member LSP taking a different path to the same destination, the ingress router is configured with all the parameters necessary to compute an RSVP ECMP LSP. These parameters along with the forwarding state information have to be synchronized between the master and backup Routing Engines to enable the support for nonstop

active routing (NSR) for container LSPs. While some of the forwarding state information on the backup Routing Engine is locally built based on the configuration, most of it is built based on periodic updates from the master Routing Engine. The container LSPs are created dynamically using the replicated states on the backup Routing Engine.

By default, normalization occurs once in every 6 hours and during this time, a number of autobandwidth adjustments happen over each member LSP. A member LSP is resized according to the traffic it carries and the configured autobandwidth configuration parameters. The aggregate demand on a container LSP is tracked by summing up the bandwidth across all the member LSPs.

For RSVP point-to-point LSPs, a Routing Engine switchover can be under any one of the following:

- **Steady state**

In the steady state, the LSP state is up and forwards traffic; however, no other event, such as the make-before-break (MBB), occurs on the LSP. At this stage, the RPD runs on both the Routing Engines, and the switchover event toggles between the master and backup Routing Engine. The backup Routing Engine has the LSP information replicated already. After the switchover, the new master uses the information of the replicated structure to construct the container LSP and en-queues the path (ERO) of LSP in the retrace mode. RSVP signals and checks if the path mentioned in the ERO is reachable. If the RSVP checks fail, then the LSP is restarted. If the RSVP checks succeed, the LSP state remains up.

- **Action leading to make-before-break (MBB)**

A container LSP can be optimized with updated bandwidth, and this change is done in a MBB fashion. During an MBB process, there are two path instances for a given LSP, and the LSP switches from one instance to another. For every Routing Engine switchover, the path is checked to find out where in the MBB process the path is. If the path is in the middle of the MBB process, with the main instance being down and the re-optimized path being up, then MBB can switch over to the new instance. The **show mpls lsp extensive** command output, in this case, is as follows:

```
13 Dec 3 01:33:38.941 Make-before-break: Switched to new instance
12 Dec 3 01:33:37.943 Record Route: 10.1.1.1
11 Dec 3 01:33:37.942 Up
10 Dec 3 01:33:37.942 Automatic Autobw adjustment succeeded: BW changes
from 100 bps to 281669 bps
9 Dec 3 01:33:37.932 Originate make-before-break call
8 Dec 3 01:33:37.931 CSPF: computation result accepted 10.1.1.1
7 Dec 3 01:28:44.228 CSPF: ERO retrace was successful 10.1.1.1
6 Dec 3 01:19:39.931 10.1.1.2 Down: mbb/reopt
5 Dec 3 01:18:29.286 Up: mbb/reopt
4 Dec 3 01:14:47.119 10.1.1.2 Down: mbb/reopt
3 Dec 3 01:13:29.285 Up: mbb/reopt
2 Dec 3 01:10:59.755 Selected as active path: selected by master RE
```

A similar behavior is retained for member LSPs during bandwidth optimization.

A Routing Engine switchover under the steady state (when normalization is not in progress), keeps the container LSPs up and running without any traffic loss. Events,

such as an MBB due to autobandwidth adjustments, link status being down, or double failure, in the steady state are similar to a normal RSVP point-to-point LSP.

If the container LSP is in the process of normalization, and the normalization event is triggered either manually or periodically, it goes through the computation and execution phase. In either of the cases, zero percent traffic loss is not guaranteed.

- Normalization in the computation phase

During the computation phase, the master Routing Engine calculates the targeted member LSP count and bandwidth with which each member LSP should be re-signaled. The backup Routing Engine has limited information about the container LSP, such as the LSP name, LSP ID, current bandwidth of its member LSP, member LSP count, and the normalization retry count. If the switchover happens during the computation phase, then the backup Routing Engine is not aware of the targeted member LSP count and the bandwidth to be signaled. Since traffic statistics are not copied to the backup Routing Engine, it cannot compute the targeted member count and bandwidth. In this case, the new master Routing Engine uses the old data stored in the targeted member LSP count and the targeted bandwidth to signal the LSPs.

- Normalization in the execution phase

During the execution phase, RSVP of the master Routing Engine tries to signal the LSPs with the newly calculated bandwidth. If the switchover occurs during the signaling of LSPs with greater bandwidth or during LSP splitting or merging, then the new master Routing Engine uses the information of the targeted member count and bandwidth value to be signaled with, to bring up the LSPs.

### ***IPG-FA Support***

A forwarding adjacency (FA) is a traffic engineering label-switched path (LSP) that is configured between two nodes and used by an interior gateway protocol (IGP) to forward traffic. By default, an IGP does not consider MPLS traffic-engineering tunnels between sites, for traffic forwarding. Forwarding adjacency treats a traffic engineering LSP tunnel as a link in an IGP topology, thus allowing the nodes in the network also to forward the IP traffic to reach the destination over this FA LSP. A forwarding adjacency can be created between routing devices regardless of their location in the network.

To advertise a container LSP as an IGP-FA, the LSP name needs to be configured either under IS-IS or OSPF. For example:

```
IS-IS  [edit]
       protocols {
         isis {
           label-switched-path container-lsp-name;
         }
       }

OSPF   [edit]
       protocols {
         ospf {
           area 0.0.0.0 {
             label-switched-path container-lsp-name;
           }
         }
       }
```

```

    }
  }

```



**NOTE:** The IGP-FA is applied to both container LSPs and regular point-to-point LSPs. If a container LSP and a point-to-point LSP share the same name, the point-to-point LSP is given preference for FA.

### Static Route Support

Static routes often include only one or very few paths to a destination and generally do not change. These routes are used for stitching services when policies and other protocols are not configured.

To advertise a container LSP as a static route, the LSP name needs to be configured under the static route configuration. For example:

```

Static Route [edit]
routing-options {
  static {
    route destination {
      lsp-next-hop container-lsp-name;
    }
  }
}

```



**NOTE:** The static route support is applied to both container LSPs and regular point-to-point LSPs. If a container LSP and a point-to-point LSP share the same name, the point-to-point LSP is given preference for static routing.

## Configuration Statements Supported for Container LSPs

Table 64 on page 504 lists the MPLS LSP configuration statements that apply to RSVP LSP and a container LSP (nominal and supplementary).

The configuration support is defined using the following terms:

- Yes—The configuration statement is supported for this type of LSP.
- No—The configuration statement is not supported for this type of LSP.
- N/A—The configuration statement is not applicable for this type of LSP.

**Table 64: Applicability of RSVP LSPs Configuration to a Container LSP**

Configuration Statement	RSVP LSP (Ingress)	Member LSP (Ingress)
adaptive	Yes	Yes
(Default: non-adaptive)		



Table 64: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

Configuration Statement	RSVP LSP (Ingress)	Member LSP (Ingress)
admin-down	Yes	Yes
admin-group	Yes	Yes
admin-groups-except	Yes	Yes
apply-groups	Yes	Yes
apply-groups-except	Yes	Yes
associate-backup-pe-groups	Yes	No
associate-lsp (No bidirectional support)	Yes	No
auto-bandwidth	Yes	Yes
backup	Yes	No
bandwidth	Yes	Yes
class-of-service	Yes	Yes
corouted-bidirectional (No bidirectional support)	Yes	No
corouted-bidirectional-passive (No bidirectional support)	Yes	No
description	Yes	Yes
disable	Yes	Yes
egress-protection	Yes	No
exclude-srlg	Yes	Yes
fast-reroute (Same fast reroute for all member LSPs)	Yes	Yes
from	Yes	Yes
hop-limit	Yes	Yes

Table 64: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

Configuration Statement	RSVP LSP (Ingress)	Member LSP (Ingress)
install	Yes	Yes
inter-domain (Same termination router)	Yes	Yes
secondary (All LSPs are primary)	Yes	No
ldp-tunneling (All LSPs do tunneling)	Yes	Yes
least-fill	Yes	Yes
link-protection (All LSPs share same link protection mechanism)	Yes	Yes
lsp-attributes	Yes	Yes
lsp-external-controller	Yes	No
metric (All LSPs are same)	Yes	Yes
most-fill	Yes	Yes
no-cspf (LSPs use IGP)	Yes	Yes
no-decrement-ttl (All LSPs share same TTL behavior)	Yes	Yes
no-install-to-address	Yes	Yes
no-record	Yes	Yes
node-link-protection (All LSPs share same node-link protection mechanism)	Yes	Yes
oam	Yes	Yes

Table 64: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

Configuration Statement	RSVP LSP (Ingress)	Member LSP (Ingress)
optimize-hold-dead-delay (All LSPs have same value)	Yes	Yes
optimize-switchover-delay (All LSPs have same value)	Yes	Yes
optimize-timer (All LSPs have same value)	Yes	Yes
p2mp	Yes	N/A
policing (Variable traffic)	Yes	No
preference	Yes	Yes
primary (All paths are primary)	Yes	No
random	Yes	Yes
record	Yes	Yes
retry-limit (Applicable to members)	Yes	Yes
retry-timer (Applicable to members)	Yes	Yes
revert-timer (No secondary LSP)	Yes	No
secondary (All LSPs are primary)	Yes	No
soft-preemption	Yes	Yes
standby (All LSPs are standby)	Yes	No

Table 64: Applicability of RSVP LSPs Configuration to a Container LSP (*continued*)

Configuration Statement	RSVP LSP (Ingress)	Member LSP (Ingress)
template	Yes	No
to	Yes	Yes
traceoptions	Yes	Yes
ultimate-hop-popping	Yes	Yes

### Impact of Configuring Container LSPs on Network Performance

A container LSP is a container LSP that allows multiple member LSPs to co-exist and be managed as a bundle. The member LSPs are similar to independent point-to-point RSVP LSPs. As a result, resource consumption is similar to the sum of resources consumed by each point-to-point RSVP LSP. However, provisioning a container LSP is more efficient, as under-utilized member LSPs are dynamically removed, thus saving memory and CPU resources.

The container LSP features are dependent on the presence of a functional base MPLS RSVP implementation. As a result, a container LSP does not introduce any security considerations beyond the existing considerations for the base MPLS RSVP functionality. The categories of possible attacks and countermeasures are as follows:

- Interaction with processes and router configuration

No new communication mechanisms with external hosts are required for a container LSP. Data arrives at the RSVP module through local software processes and router configuration, other than RSVP neighbor adjacency. Junos OS provides security controls on access to the router and router configuration.

- Communication with external RSVP neighbors

RSVP signaled MPLS LSPs depend on the services of RSVP and IGP to communicate RSVP messages among neighboring routers across the network. Because the RSVP sessions involve communication outside of the local router, they are subject to many forms of attack, such as spoofing of peers, injection of falsified RSVP messages and route updates, and attacks on the underlying TCP/UDP transport for sessions. Junos OS provides countermeasures for such attack vectors.

- Resource limits and denial of service

Junos OS provides several mechanisms through policers and filters to protect against denial-of-service attacks based on injecting higher than the expected traffic demands. At the MPLS LSP level, Junos OS allows operators to configure limits on the LSP bandwidth and the number of LSPs. However, like point-to-point RSVP LSPs, container LSPs do not enforce limits on the volume of traffic forwarded over these LSPs.

## Supported and Unsupported Features

Junos OS supports the following container LSP features:

- Equal-bandwidth-based LSP splitting mechanism
- Aggregate-bandwidth-based LSP splitting and merging in a make-before-break way
- LSP-number-based naming mechanism for dynamically created member LSPs
- Periodic sampling mechanisms to estimate aggregate bandwidth
- Interoperability with auto-bandwidth feature
- ECMP using the dynamically created LSPs
- LDP-tunneling on the dynamically created LSP
- Configuring container LSP using IGP shortcuts
- Aggregated Ethernet links
- Logical systems

Junos OS does **not** support the following container LSP functionality:

- Node and link disjoint paths for different LSPs between an ingress and an egress routing device
- Bandwidth allocation policy different from equal bandwidth policy at the normalization event
- Constraint-based routing path computation to find equal IGP cost paths for different LSPs
- RSVP objects, such as **MLSP\_TUNNEL Sender Template**, and **MLSP\_TUNNEL Filter Specification** defined in [KOMPELLA-MLSP]
- Change in topology as a trigger for LSP splitting and merging
- Change in topology and link failure as a trigger for normalization, unless member LSPs go down
- Egress protection on container LSP
- Container LSP as a backup LSP for IGP interface
- Container LSP configured as IGP interface as forwarding address
- Container LSP as provider tunnel for multicast VPNs
- Dynamic LSPs for normalization
- CCC using container LSP
- Secondary paths for container LSP
- Bidirectional container LSP
- Policing
- Static routes using container LSPs as next hops on a best-effort basis

- External path computing entity, such as PCE
- Graceful Routing Engine switchover
- Nonstop active routing
- Unified ISSU
- Multichassis
- IPv6

**Related  
Documentation**

- *Example: Configuring Dynamic Bandwidth Management Using Container LSP*
- [Maximize Bandwidth Utilization with Juniper Networks TE++](#)

## CHAPTER 8

# Configuration Statements for RSVP

- [admin-group](#) on page 513
- [authentication-key \(Protocols RSVP\)](#) on page 514
- [aggregate \(Protocols RSVP\)](#) on page 515
- [bandwidth \(Protocols RSVP\)](#) on page 516
- [bypass \(Signaled LSP\)](#) on page 517
- [class-of-service \(Protocols RSVP\)](#) on page 518
- [container-label-switched-path](#) on page 519
- [disable \(Protocols RSVP\)](#) on page 520
- [fast-reroute \(Protocols RSVP\)](#) on page 521
- [graceful-deletion-timeout](#) on page 521
- [graceful-restart \(Enabling Globally\)](#) on page 522
- [hello-acknowledgements](#) on page 523
- [hello-interval \(Protocols RSVP\)](#) on page 524
- [helper-disable \(Multiple Protocols\)](#) on page 525
- [hop-limit](#) on page 526
- [interface \(Protocols RSVP\)](#) on page 527
- [keep-multiplier](#) on page 528
- [link-protection \(RSVP\)](#) on page 529
- [load-balance \(Protocols RSVP\)](#) on page 530
- [max-bypasses](#) on page 531
- [maximum-helper-recovery-time](#) on page 532
- [maximum-helper-restart-time \(RSVP\)](#) on page 533
- [no-cspf \(Protocols RSVP\)](#) on page 534
- [no-interface-hello](#) on page 534
- [no-local-reversion](#) on page 535
- [no-node-id-subobject](#) on page 536
- [no-p2mp-sublsp](#) on page 536
- [node-hello](#) on page 537

- [optimize-timer \(Protocols RSVP\) on page 537](#)
- [path \(Protocols RSVP\) on page 538](#)
- [preemption on page 539](#)
- [priority \(Protocols RSVP\) on page 540](#)
- [refresh-time on page 541](#)
- [reliable on page 541](#)
- [setup-protection on page 542](#)
- [subscription on page 543](#)
- [soft-preemption \(Protocols RSVP\) on page 544](#)
- [splitting-merging on page 545](#)
- [traceoptions \(Protocols RSVP\) on page 547](#)
- [tunnel-services \(RSVP\) on page 549](#)
- [update-threshold on page 549](#)



## admin-group

<b>Syntax</b>	<pre>admin-group {   exclude [ <i>group-names</i> ];   include-all [ <i>group-names</i> ];   include-any [ <i>group-names</i> ]; }</pre>
<b>Hierarchy Level</b>	<pre>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</pre>
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2.
<b>Description</b>	<p>Enable you to configure administrative groups for bypass label-switched paths (LSPs). You can configure administrative groups either globally for all bypass LSPs traversing an interface or for just a specific bypass LSP.</p>
<b>Options</b>	<p><b>exclude <i>group-names</i></b>—Specify the administrative groups to exclude for a bypass LSP.</p> <p><b>include-all <i>group-names</i></b>—Specify the administrative groups whose links the bypass LSP must traverse.</p> <p><b>include-any <i>group-names</i></b>—Specify the administrative groups whose links the bypass LSP can traverse.</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>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li> </ul>

## authentication-key (Protocols RSVP)

---

<b>Syntax</b>	authentication-key <i>key</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i> ], [edit protocols rsvp interface <i>interface-name</i> ], [edit protocols rsvp peer-interface <i>peer-interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	<p>Authentication key (password). Neighboring routers use the password to verify the authenticity of packets sent from this interface or peer interface.</p> <p>RSVP uses HMAC-MD5 authentication, which is defined in RFC 2104, <i>HMAC: Keyed-Hashing for Message Authentication</i>.</p> <p>All routers that are connected to the same IP subnet must use the same authentication scheme and password.</p>
<b>Options</b>	<b>key</b> —Authentication password. It can be 1 through 16 contiguous digits or letters. Separate decimal digits with periods. Separate hexadecimal digits with periods and precede the string with 0x. If you include spaces in the password, enclose the entire password in quotation marks (" ").
<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>• <i>Configuring RSVP Interfaces</i></li></ul>

## aggregate (Protocols RSVP)

<b>Syntax</b>	(aggregate   no-aggregate);
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i> ], [edit protocols rsvp interface <i>interface-name</i> ], [edit protocols rsvp peer-interface <i>peer-interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	<p>Control the use of RSVP aggregate messages on an interface or peer interface:</p> <ul style="list-style-type: none"> <li>• <b>aggregate</b>—Use RSVP aggregate messages.</li> <li>• <b>no-aggregate</b>—Do not use RSVP aggregate messages.</li> </ul> <p>Aggregate messages can pack multiple RSVP messages into a single transmission, thereby reducing network overhead and enhancing efficiency. The number of supportable sessions and processing overhead are significantly improved when aggregation is enabled.</p> <p>Not all routers connected to a subnet need to support aggregation simultaneously. Each RSVP router negotiates its intention to use aggregate messages on a per-neighbor basis. Only when both routers agree are aggregate messages sent.</p> <p>To have refresh reduction and reliable delivery, you must include the <b>aggregate</b> and <b>reliable</b> statements.</p>
<b>Default</b>	Aggregation is disabled.
<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>• <i>Configuring RSVP Interfaces</i></li> <li>• <a href="#">reliable on page 541</a></li> </ul>

## bandwidth (Protocols RSVP)

---

<b>Syntax</b>	<code>bandwidth <i>bps</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i></code> <code>link-protection],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i></code> <code>link-protection bypass <i>bypass-name</i>],</code> <code>[edit protocols rsvp interface <i>interface-name</i>],</code> <code>[edit protocols rsvp interface <i>interface-name</i> link-protection],</code> <code>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	<p>For certain logical interfaces (such as Asynchronous Transfer Mode [ATM], Permanent Virtual Circuit [PVC], or Frame Relay), you cannot determine the correct bandwidth from the hardware. This statement enables you to specify the actual available bandwidth.</p> <p>This statement also enables you to specify the bandwidth for a bypass label switched path (LSP). If you have configured multiple bypasses, this statement is mandatory and is applied to all of the bypass LSPs.</p>
<b>Default</b>	The hardware raw bandwidth is used.
<b>Options</b>	<p><b><i>bps</i></b>—Bandwidth in bits per second. You can specify this as an integer value. If you do so, count your zeros carefully, or you can use the abbreviations <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>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li><li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li><li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li></ul>

## bypass (Signaled LSP)

<b>Syntax</b>	<pre> bypass <i>bypass-name</i> {     <b>bandwidth</b> <i>bps</i>;     <b>description</b> <i>text</i>;     <b>hop-limit</b> <i>number</i>;     <b>no-cspf</b>;     <b>path</b> <i>address</i> &lt;strict   loose&gt;;     <b>priority</b> <i>setup-priority reservation-priority</i>;     <b>to address</b>; } </pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>The <b>description</b> option was added in Junos OS Release 10.4.</p>
<b>Description</b>	<p>Enables you to configure specific bandwidth and path constraints for a bypass LSP. It is possible to individually configure multiple bypass LSPs. If you do not configure the bypass LSPs individually, they all share the same path and bandwidth constraints.</p> <p>If you specify the <b>bandwidth</b>, <b>hop-limit</b>, and <b>path</b> statements for the bypass LSP, these values take precedence over the values configured at the [edit protocols rsvp interface <i>interface-name</i> link-protection] hierarchy level. The other attributes (<b>subscription</b>, <b>no-node-protection</b>, and <b>optimize-timer</b>) are inherited from the general constraints.</p>
<b>Options</b>	<p><b>bypass-name</b>—(Required) Specify a name for the bypass LSP. The name can be up to 64 characters.</p> <p><b>description</b>—Provides a textual description of the bypass LSP. Enclose any descriptive text that includes spaces in quotation marks (" "). Any descriptive text you include is displayed in the output of the <b>show mpls lsp bypass detail</b> command and has no effect on the operation of the bypass LSP. The description text can be no more than 80 characters in length.</p> <p><b>to address</b>—(Required) Specify the address for the interface of the immediate next-hop node (for link protection) or the next-next-hop node (for node-link protection). The address specified determines whether this is a link protection bypass or a node-link protection bypass. On multiaccess networks (for example, a LAN), this address is also used to specify which next-hop node is being protected.</p> <p>The remaining statements are explained separately.</p>
<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>Configuring Link Protection on Interfaces Used by LSPs</li> </ul>

## class-of-service (Protocols RSVP)

---

<b>Syntax</b>	<code>class-of-service <i>cos-value</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</code> <code>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</code> <code>[edit protocols rsvp interface <i>interface-name</i> link-protection],</code> <code>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</code>
<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 bypass LSP. You can specify a single CoS value for all the bypass LSPs traversing an interface. You can also configure CoS values for specific bypass LSPs traversing an interface.</p> <p>The CoS value might affect the scheduling or queuing algorithm of traffic traveling along an LSP.</p>
<b>Options</b>	<p><b><i>cos-value</i></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>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li></ul>

## container-label-switched-path

<b>Syntax</b>	<pre> container-label-switched-path <i>lsp-name</i> {   disable;   description <i>description</i>;   label-switched-path-template;   splitting-merging;   suffix <i>string</i>;   to <i>ip-address</i>; }</pre>
<b>Hierarchy Level</b>	[edit protocols mpls]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 14.2.</p> <p>Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.</p>
<b>Description</b>	Configure a multi-label-switched path (LSP) tunnel between the ingress and the egress routers. The container LSP consists of several member LSPs to the same destination.
<b>Options</b>	<p><b>disable</b>—Disable MPLS container-label-switched path.</p> <p><b>description <i>description</i></b>—Text describing the container LSP.</p> <p><b>suffix <i>string</i></b>—Suffix to generate names of member LSPs of the container LSP.</p> <p><b>to <i>ip-address</i></b>—IP address of the egress 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>

## disable (Protocols RSVP)

---

<b>Syntax</b>	disable;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit logical-systems <i>logical-system-name</i> protocols rsvp graceful-restart], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i> ], [edit protocols rsvp], [edit protocols rsvp graceful-restart], [edit protocols rsvp interface <i>interface-name</i> ], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp peer-interface <i>peer-interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Explicitly disable RSVP or RSVP graceful restart. Explicitly disable link protection on the specified interface.
<b>Default</b>	RSVP is enabled on interfaces and peer interfaces configured with the RSVP <b>interface</b> statement. RSVP graceful restart is enabled on the router. 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>• <i>Minimum RSVP Configuration</i></li><li>• <i>Configuring RSVP Graceful Restart</i></li><li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li></ul>



## fast-reroute (Protocols RSVP)

<b>Syntax</b>	<code>fast-reroute optimize-timer <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
<b>Release Information</b>	Statement added in Junos OS Release 7.5. Statement introduced in Junos OS Release 14.1 for the QFX Series.
<b>Description</b>	Configure the optimize timer for fast reroute. The optimize timer triggers a periodic optimization process that recomputes the fast reroute detour LSPs to use network resources more efficiently.
<b>Options</b>	<b><i>seconds</i></b> —Specify the number of seconds between fast reroute detour LSP optimizations. <b>Range:</b> 0 through 65,535 seconds <b>Default:</b> 0 (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>• <i>Configuring the Optimization Interval for Fast Reroute Paths</i></li> </ul>

## graceful-deletion-timeout

<b>Syntax</b>	<code>graceful-deletion-timeout <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the time, in seconds, before completing graceful deletion of signaling.
<b>Options</b>	<b><i>seconds</i></b> —Time before completing graceful deletion of signaling. <b>Range:</b> 1 through 300 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>• <i>Gracefully Tearing Down GMPLS LSPs</i></li> </ul>

## graceful-restart (Enabling Globally)

<b>Syntax</b>	<pre> graceful-restart {   disable;   helper-disable;   maximum-helper-recovery-time seconds;   maximum-helper-restart-time seconds;   notify-duration seconds;   recovery-time seconds;   restart-duration seconds;   stale-routes-time seconds; } </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 9.0 for EX Series switches. Statement introduced in Junos OS Release 12.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.



### NOTE:

- For VPNs, the **graceful-restart** statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.
- For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.
- LDP sessions flap when **graceful-restart** configurations change.

<b>Default</b>	Graceful restart is disabled by default.
<b>Options</b>	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>• <i>Enabling Graceful Restart</i></li> <li>• <i>Configuring Routing Protocols Graceful Restart</i></li> </ul>

- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- *Graceful Restart Configuration Statements*
- *Configuring Graceful Restart for QFabric Systems*

## hello-acknowledgements

---

<b>Syntax</b>	hello-acknowledgements;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-systems-name</i> protocols rsvp], [edit protocols rsvp]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable hello messages from nonsession neighbors to be acknowledged with a hello acknowledgment message. Once hello acknowledgments are enabled, the router continues to acknowledge hello messages from any nonsession RSVP neighbors unless the interface itself goes down or the configuration is changed by an administrator.
<b>Default</b>	Hello acknowledgments are 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>• <i>Configuring Hello Acknowledgments for Nonsession RSVP Neighbors</i></li> </ul>

## hello-interval (Protocols RSVP)

---

<b>Syntax</b>	hello-interval <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i> ], [edit protocols rsvp interface <i>interface-name</i> ], [edit protocols rsvp peer-interface <i>peer-interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable the sending of hello packets on the interface.
<b>Options</b>	<b><i>seconds</i></b> —Length of time between hello packets. A value of 0 disables the sending of hello packets on the interface. <b>Range:</b> 1 through 60 seconds <b>Default:</b> 9 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>• <i>Configuring RSVP Interfaces</i></li></ul>

## helper-disable (Multiple Protocols)

<b>Syntax</b>	helper-disable;
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols (isis   ldp   ospf   ospf3   rsvp) <a href="#">graceful-restart</a>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols (ldp   ospf   ospf3) <a href="#">graceful-restart</a>],</p> <p>[edit protocols (isis   ldp   ospf   ospf3   rsvp) <a href="#">graceful-restart</a>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols (ldp   ospf   ospf3) <a href="#">graceful-restart</a>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	Disable helper mode for graceful restart. When helper mode is disabled, a router or switch cannot help a neighboring router that is attempting to restart.
<b>Default</b>	Helper mode is enabled by default for these supported protocols: IS-IS, LDP, OSPF/OSPFv3, and 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>• <i>Configuring Routing Protocols Graceful Restart</i></li> <li>• <i>Configuring Graceful Restart for MPLS-Related Protocols</i></li> </ul>

## hop-limit

<b>Syntax</b>	<code>hop-limit <i>number</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 <a href="#">label-switched-path</a> <i>lsp-name</i>],  [edit logical-systems <i>logical-system-name</i> protocols mpls label-switched-path <i>lsp-name</i> <a href="#">fast-reroute</a>],  [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 logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],  [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],  [edit protocols mpls],  [edit protocols mpls <a href="#">label-switched-path</a> <i>lsp-name</i>],  [edit 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="#">primary</a>   <a href="#">secondary</a>) <i>path-name</i>],  [edit protocols rsvp interface <i>interface-name</i> link-protection],  [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	<p>Specify the maximum number of routers that an LSP can traverse. This limit can be applied to any of the following:</p> <ul style="list-style-type: none"> <li>LSPs—The configured hop limit includes the ingress and egress routers. You can specify a hop limit for an LSP and for both primary and secondary paths.</li> <li>Fast reroute detour—Specify the number of additional routers a fast reroute detour can traverse relative to the protected LSP. For example, if an LSP traverses 4 routers, any detour for the LSP can be no more than 10 router hops, including the ingress and egress routers.</li> <li>Link protection bypass—Specify the maximum number of routers that a link protection bypass can traverse.</li> </ul>
<b>Options</b>	<p><b><i>number</i></b>—Maximum number of hops.</p> <p><b>Range:</b> 2 through 255 (for an LSP or for a link protection bypass); 0 through 255 (for fast reroute)</p> <p><b>Default:</b> 255 (for an LSP or for a link protection bypass); 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</a></li> <li><a href="#">Limiting the Number of Hops in LSPs</a></li> <li><a href="#">Configuring Link Protection on Interfaces Used by LSPs</a></li> </ul>

## interface (Protocols RSVP)

```

Syntax  interface interface-name {
        disable;
        (aggregate | no-aggregate);
        authentication-key key;
        bandwidth bps;
        hello-interval seconds;
        link-protection {
            disable;
            admin-group {
                exclude [ group-names ];
                include-all [ group-names ];
                include-any [ group-names ];
            }
            bandwidth bps;
            bypass bypass-name {
                bandwidth bps {
                    ct0 bps;
                    ct1 bps;
                    ct2 bps;
                    ct3 bps;
                }
                description text;
                class-of-service cos-value;
                hop-limit number;
                no-cspf;
                path address <strict | loose>;
                priority setup-priority reservation-priority;
                to address;
            }
            class-of-service cos-value;
            hop-limit number;
            max-bypasses number;
            no-cspf;
            no-node-protection;
            optimize-timer seconds;
            path address <strict | loose>;
            priority setup-priority reservation-priority;
            subscription percentage;
        }
        (reliable | no-reliable);
        subscription percentage {
            ct0 percentage;
            ct1 percentage;
            ct2 percentage;
            ct3 percentage;
        }
        update-threshold threshold;
    }

```

**Hierarchy Level** [edit logical-systems *logical-system-name* protocols rsvp],  
[edit protocols rsvp]

<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable RSVP on one or more router interfaces.
<b>Default</b>	RSVP is disabled on all interfaces.
<b>Options</b>	<b><i>interface-name</i></b> —Name of an interface. To configure all interfaces, specify <b>all</b> . For details about specifying interfaces, see the <i>Junos OS Network Interfaces Library for Routing Devices</i> .  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>• <i>Minimum RSVP Configuration</i></li></ul>

---

## keep-multiplier

---

<b>Syntax</b>	keep-multiplier <i>number</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Set the keep multiplier value.
<b>Options</b>	<b><i>number</i></b> —Multiplier value. <b>Range:</b> 1 through 255 <b>Default:</b> 3
<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>• <i>Configuring Timers for RSVP Refresh Messages</i></li></ul>



## link-protection (RSVP)

<b>Syntax</b>	<pre> link-protection {   disable;   admin-group {     exclude [ group-names ];     include-all [ group-names ];     include-any [ group-names ];   }   bandwidth bps;   bypass bypass-name {     bandwidth bps {       ct0 bps;       ct1 bps;       ct2 bps;       ct3 bps;     }     description text;     class-of-service cos-value;     hop-limit number;     no-cspf;     path address &lt;strict   loose&gt;;     priority setup-priority reservation-priority;     to address;   }   class-of-service cos-value;   hop-limit number;   max-bypasses number;   no-cspf;   no-node-protection;   optimize-timer seconds;   path address &lt;strict   loose&gt;;   priority setup-priority reservation-priority;   subscription percentage; } </pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ], [edit protocols rsvp interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 14.1X53-D10 for the QFX Series and for EX4600 switches.
<b>Description</b>	Enable link protection on the specified interface. Using link protection, you can configure a network to reroute traffic quickly around broken links. To fully enable link protection, you also need to configure the <b>link-protection</b> statement at the [edit protocols mpls label-switched-path <i>lsp-name</i> ] hierarchy level. You can configure single or multiple bypasses for protected interface.
<b>Default</b>	Link protection is disabled.

**Options**    **no-node-protection**—Disable node-link protection on the RSVP interface. Link protection remains active. When this option is configured, the router can only initiate a next-hop bypass, not a next-next-hop bypass.

The remaining statements are explained separately.

**Required Privilege Level**    routing—To view this statement in the configuration.  
   routing-control—To add this statement to the configuration.

**Related Documentation**    • *Configuring Link Protection on Interfaces Used by LSPs*  
   • *link-protection (Dynamic LSPs)*

---

## load-balance (Protocols RSVP)

---

**Syntax**    load-balance {  
                 bandwidth;  
                 }

**Hierarchy Level**    [edit logical-systems *logical-system-name* protocols rsvp],  
                                 [edit protocols rsvp]

**Release Information**    Statement introduced before Junos OS Release 7.4.  
                                 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description**    Load-balance traffic between RSVP LSPs.

**Options**    **bandwidth**—Load-balance traffic between RSVP LSPs based on the bandwidth configured for each LSP.

**Required Privilege Level**    routing—To view this statement in the configuration.  
   routing-control—To add this statement to the configuration.

**Related Documentation**    • *Configuring Load Balancing Across RSVP LSPs*

## max-bypasses

<b>Syntax</b>	<code>max-bypasses <i>number</i>;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ], [edit protocols rsvp interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Range modified in Junos OS Release 9.3.
<b>Description</b>	Specify the maximum number of dynamic bypass LSPs permitted for protecting this interface. When this option is configured, multiple bypasses for link protection are enabled. Call admission control (CAC) is also enabled. The limit on bypasses configured applies only to dynamically generated bypass LSPs. By default, this option is disabled and only one dynamic bypass LSP is enabled for each interface. If you configure <b>max-bypasses</b> , you must also configure the <b>bandwidth</b> statement.
<b>Options</b>	<b>number</b> —Configure the maximum number of bypass LSPs. If you configure a value of 0, no dynamic bypass LSPs are allowed to be established for the interface. Only static bypass LSPs can be configured. <b>Range:</b> 0 through 99 <b>Default:</b> 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>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li> </ul>

## maximum-helper-recovery-time

---

<b>Syntax</b>	<code>maximum-helper-recovery-time seconds;</code>
<b>Hierarchy Level</b>	[edit protocols rsvp <a href="#">graceful-restart</a> ], [edit logical-systems <i>logical-system-name</i> protocols rsvp <a href="#">graceful-restart</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the length of time the router or switch retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart.
<b>Options</b>	<b>seconds</b> —Length of time that the router retains the state of its Resource Reservation Protocol (RSVP) neighbors while they undergo a graceful restart. <b>Range:</b> 1 through 3600 <b>Default:</b> 180
<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>• <i>Configuring Graceful Restart Options for RSVP, CCC, and TCC</i></li><li>• <a href="#">maximum-helper-restart-time (RSVP) on page 533</a></li></ul>

## maximum-helper-restart-time (RSVP)

<b>Syntax</b>	<code>maximum-helper-restart-time <i>seconds</i>;</code>
<b>Hierarchy Level</b>	[edit protocols rsvp <a href="#">graceful-restart</a> ], [edit logical-systems <i>logical-system-name</i> protocols rsvp <a href="#">graceful-restart</a> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Specify the length of time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down. This value is applied to all RSVP neighbor routers and should be based on the time that the slowest RSVP neighbor requires for restart.
<b>Options</b>	<b><i>seconds</i></b> —The time the router or switch waits after it discovers that a neighboring router has gone down before it declares the neighbor down. <b>Range:</b> 1 through 1800 <b>Default:</b> 60
<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><i>Configuring Graceful Restart Options for RSVP, CCC, and TCC</i></li> <li><a href="#">maximum-helper-recovery-time on page 532</a></li> </ul>

## no-cspf (Protocols RSVP)

<b>Syntax</b>	no-cspf;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection], [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i> ], [edit protocols rsvp interface <i>interface-name</i> link-protection], [edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 7.5.
<b>Description</b>	Disable CSPF computation on all bypass LSPs or on a specific bypass LSP. You need to disable CSPF for link protection to function properly on interarea paths.
<b>Default</b>	CSPF is 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>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li> </ul>

## no-interface-hello

<b>Syntax</b>	no-interface-hello;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
<b>Release Information</b>	Statement introduced in JUNOS Release 10.0. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Allows you to explicitly disable RSVP interface hellos globally on the router. This type of configuration might be necessary in networks where the Juniper Networks router has numerous RSVP connections with equipment from other vendors. However, if you disable RSVP interface hellos globally, you can also configure a hello interval on an RSVP interface using the <a href="#">hello-interval (Protocols RSVP)</a> statement. This configuration disables RSVP interface hellos globally but enables RSVP interface hellos on the specified interface. This configuration might be necessary in a heterogeneous network where some devices support RSVP node ID hellos and other devices support RSVP interface hellos.
<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>• <i>Configuring RSVP Node ID Hellos</i></li> <li>• <a href="#">hello-interval (Protocols RSVP) on page 524</a></li> </ul>

## no-local-reversion

<b>Syntax</b>	no-local-reversion;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.4.
<b>Description</b>	<p>Disables RSVP local revertive mode as specified in RFC 4090, <i>Fast Reroute Extensions to RSVP-TE for LSP Tunnels</i>. RSVP local revertive mode is supported on all Juniper Networks routers running the Junos OS. It is the default behavior. If you include this statement, the Juniper Networks router uses global revertive mode instead. You might need to disable RSVP local revertive mode on Juniper Networks routers if your network includes equipment that does not support this mode.</p> <p>The following information can also be found in RFC 4090. Refer to the full RFC for additional information. When an LSP fails, the connection can be repaired locally using a traffic protection mechanism such as fast reroute. To restore the LSP to a full working path, RFC 4090 specifies the following strategies:</p> <ul style="list-style-type: none"> <li>• Local revertive mode—Upon detecting that the path is restored, the point of local repair (PLR) resignals each of the LSPs that were formerly routed over the restored path. Every LSP successfully resignaled along the restored path is switched back.</li> <li>• Global revertive mode—The ingress router of each tunnel is responsible for reoptimizing the LSPs that used the failed path. There are several potential reoptimization triggers: RSVP error messages, inspection of OSPF LSAs or IS-IS LSPs, and timers. This re-optimization process can proceed as soon as the failure is detected. It is not tied to the restoration of the failed path.</li> </ul>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## no-node-id-subobject

---

<b>Syntax</b>	no-node-id-subobject;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Disable the record route object (RRO) node ID subobject for compatibility with earlier versions of the Junos OS. To interoperate with other vendors' equipment, the Junos OS supports the RRO node ID subobject for use in inter-AS link and node protection configurations.
<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>• <i>Configuring Inter-AS Node and Link Protection</i></li></ul>

## no-p2mp-sublsp

---

<b>Syntax</b>	no-p2mp-sublsp;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Reject Resv messages that include the S2L_SUB_LSP object. By default, Resv messages that include the S2L_SUB_LSP object are accepted. However, in a network which includes Juniper Networks devices running both Junos OS Release 9.2 and later and Junos OS Release 9.1 and earlier, it is necessary to configure the <b>no-p2mp-sublsp</b> statement on devices running Junos OS Release 9.2 and later to ensure that point-to-multipoint LSPs function properly.
<b>Default</b>	Resv messages that include the S2L_SUB_LSP object are accepted.
<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>• <i>Preserving Point-to-Multipoint LSP Functioning with Different Junos OS Releases</i></li></ul>



## node-hello

<b>Syntax</b>	node-hello;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
<b>Release Information</b>	Statement introduced in JUNOS Release 10.0. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enables node-ID based RSVP hellos globally on all of the RSVP interfaces on the router to allow Juniper Networks routers to interoperate with the equipment of other vendors. By default, the JUNOS Software uses interface-based RSVP hellos and node-ID based RSVP hellos are disabled. If you have not enabled RSVP node IDs on the router, the JUNOS software does not accept any node-ID hello packets.
<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>• <i>Configuring RSVP Node ID Hellos</i></li> </ul>

## optimize-timer (Protocols RSVP)

<b>Syntax</b>	optimize-timer <i>seconds</i> ;
<b>Hierarchy Level</b>	[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]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure an optimize timer for a bypass LSP. The optimize timer initiates a periodic optimization process that reshuffles data LSPs among bypass LSPs to achieve the most efficient use of network resources. The optimization process attempts to either minimize the number of bypasses currently in use, minimize the total amount of bandwidth reserved for all bypasses, or both.
<b>Options</b>	<p><b><i>seconds</i></b>—Specify the number of seconds between optimizations.</p> <p><b>Range:</b> 0 through 65,535 seconds</p> <p><b>Default:</b> 0 (disabled)</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>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li> </ul>

## path (Protocols RSVP)

<b>Syntax</b>	<code>path address &lt;strict   loose&gt;;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>]</p>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure an explicit path (a sequence of strict or loose routes) to control where and how a bypass LSP is established. If multiple bypasses are configured, they all will use the same explicit path.
<b>Default</b>	No path is configured. CSPF automatically calculates the path the bypass LSP takes.
<b>Options</b>	<p><b>address</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 bypass LSP.</p> <p><b>loose</b>—(Optional) The next address in the <b>path</b> statement is loose. The LSP can traverse other routers before reaching this router.</p> <p><b>Default:</b> <b>strict</b></p> <p><b>strict</b>—(Optional) 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>Configuring Link Protection on Interfaces Used by LSPs</li> </ul>

## preemption

<b>Syntax</b>	<pre>preemption {   (aggressive   disabled   normal);   soft-preemption {     cleanup-timer <i>seconds</i>;   } }</pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Control RSVP session preemption.
<b>Default</b>	<b>normal</b>
<b>Options</b>	<p><b>aggressive</b>—Preempt RSVP sessions whenever bandwidth is insufficient to handle all sessions. A session is preempted whenever bandwidth is lowered or a new higher-priority session is established.</p> <p><b>disabled</b>—Do not preempt RSVP sessions.</p> <p><b>normal</b>—Preempt RSVP sessions to accommodate new higher-priority sessions when bandwidth is insufficient to handle all sessions.</p> <p>The remaining statements are explained separately.</p>
<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>• <i>Preempting RSVP Sessions</i></li> </ul>

## priority (Protocols RSVP)

<b>Syntax</b>	<code>priority setup-priority reservation-priority;</code>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection],</p> <p>[edit protocols rsvp interface <i>interface-name</i> link-protection bypass <i>bypass-name</i>],</p>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the setup priority and reservation priority for a bypass LSP. If insufficient link bandwidth is available during session establishment, the setup priority is compared with other setup priorities for established sessions on the link to determine whether some of them should be preempted to accommodate the new session. The session with the lower-hold priority is preempted.
<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>	<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>Configuring Link Protection on Interfaces Used by LSPs</li> <li>Configuring Priority and Preemption for LSPs</li> </ul>

## refresh-time

<b>Syntax</b>	<code>refresh-time seconds;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Set the refresh time.
<b>Options</b>	<b>seconds</b> —Refresh time. <b>Range:</b> 1 through 65,535 <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>• <i>Configuring Timers for RSVP Refresh Messages</i></li> </ul>

## reliable

<b>Syntax</b>	<code>(reliable   no-reliable);</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ], [edit logical-systems <i>logical-system-name</i> protocols rsvp peer-interface <i>peer-interface-name</i> ], [edit protocols rsvp interface <i>interface-name</i> ], [edit protocols rsvp peer-interface <i>peer-interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable reliable message delivery on the interface.  In order to have refresh reduction and reliable delivery, you must include the <b>aggregate</b> and <b>reliable</b> statements.
<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>• <i>Configuring RSVP Interfaces</i></li> <li>• <a href="#">aggregate on page 515</a></li> </ul>

## setup-protection

---

<b>Syntax</b>	setup-protection;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
<b>Description</b>	The facility-backup fast reroute mechanism can provide setup protection for LSPs which are in the process of being signaled. Both point-to-point LSPs and point-to-multipoint LSPs are supported. You should configure the <b>setup-protection</b> statement on each of the routers along the LSP path on which you want to enable LSP setup protection. You should also configure IGP traffic engineering on all of the routers on the LSP path.
<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>• <i>Configuring RSVP Setup Protection</i></li></ul>

## subscription

<b>Syntax</b>	<pre>subscription <i>percentage</i> {     ct0 <i>percentage</i>;     ct1 <i>percentage</i>;     ct2 <i>percentage</i>;     ct3 <i>percentage</i>; }</pre>
<b>Hierarchy Level</b>	<p>[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i>],          [edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> link-protection],          [edit protocols rsvp interface <i>interface-name</i>],          [edit protocols rsvp interface <i>interface-name</i> link-protection]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.          Statement introduced in Junos OS Release 12.3X50 for the QFX Series.</p>
<b>Description</b>	<p>Configure the amount of bandwidth subscribed to a class type (when you have enabled Differentiated Services) or bypass LSP (when you have enabled link protection). <b>subscription</b> is the percentage of the link bandwidth that can be used for the RSVP reservation process.</p>
<b>Options</b>	<p><b>ctnumber percentage</b>—Percentage of the class-type bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type. You can specify bandwidth subscriptions for class types 0 through 3. This option is not available for bypass LSPs.</p> <p><b>Range:</b> 0 through 65,000  <b>Default:</b> 100 percent</p> <p><b>percentage</b>—Percentage of the class-type or bypass LSP bandwidth allowed for reservations. If you specify a value greater than 100, you are oversubscribing the class type or bypass LSP.</p> <p><b>Range:</b> 0 through 65,000  <b>Default:</b> 100 percent</p>
<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>• <i>Configuring the Bandwidth Subscription Percentage for LSPs</i></li> <li>• <i>Configuring Link Protection on Interfaces Used by LSPs</i></li> </ul>

## soft-preemption (Protocols RSVP)

---

<b>Syntax</b>	<code>soft-preemption {     cleanup-timer <i>seconds</i>; }</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp preemption], [edit protocols rsvp preemption]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable soft preemption to attempt to establish a new path for a preempted LSP before tearing it down.
<b>Options</b>	<b>cleanup-timer</b> —A value of 0 disables soft preemption. <b>Range:</b> 0 through 180 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>• <i>Configuring MPLS Soft Preemption</i></li></ul>



## splitting-merging

<b>Syntax</b>	<pre> splitting-merging {   maximum-member-lsps <i>number</i>;   maximum-signaling-bandwidth <i>bps</i>;   merging-bandwidth <i>bps</i>;   minimum-member-lsps <i>number</i>;   minimum-signaling-bandwidth <i>bps</i>;   normalization;   sampling;   splitting-bandwidth <i>bps</i>;   splitting-merging-threshold <i>percent</i>; } </pre>
<b>Hierarchy Level</b>	[edit protocols mpls container-label-switched-path <i>lsp-name</i> ]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 14.2.</p> <p>Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.</p>
<b>Description</b>	Perform splitting and merging.
<b>Options</b>	<p><b>maximum-member-lsps <i>number</i></b>—Number of label-switched paths (LSPs) that a container LSP can have as member LSPs at maximum.</p> <p><b>Default:</b> 1</p> <p><b>maximum-signaling-bandwidth <i>bandwidth</i></b>—Amount of bandwidth in bits per second (bps) that can be signaled for an LSP at maximum after normalization. When <b>maximum-signaling-bandwidth</b> is not configured, the value is derived from the <b>splitting-bandwidth</b>.</p> <p>When auto-bandwidth adjustment is done between two normalization events, per LSP auto-bandwidth configuration and thresholds are used instead of the <b>splitting-bandwidth</b>.</p> <p><b>Default:</b> 1 bps</p> <p><b>merging-bandwidth <i>bandwidth</i></b>—Amount of bandwidth in bits per second (bps) that is used for merging during normalization.</p> <p><b>Default:</b> 1 bps</p> <p><b>minimum-member-lsps <i>number</i></b>—Number of LSPd that a container LSP can have as member LSPs at minimum.</p> <p><b>Default:</b> 64</p> <p><b>minimum-signaling-bandwidth <i>bandwidth</i></b>—Amount of bandwidth in bits per second (bps) that can be signaled for an LSP at minimum after normalization. When <b>minimum-signaling-bandwidth</b> is not configured, the value is derived from the <b>merging-bandwidth</b>.</p> <p>When auto-bandwidth adjustment is done between two normalization events, per LSP auto-bandwidth configuration and thresholds are used instead of the <b>merging-bandwidth</b>.</p> <p><b>Default:</b> 1 bps</p>

**splitting-bandwidth *bandwidth***—Amount of bandwidth in bits per second (bps) that can be used for splitting during normalization.

**Default:** 1 bps

**splitting-merging-threshold *percent***—Percentage changes in aggregate bandwidth relevant for splitting and merging.

**Default:** 0%

The remaining statements are explained separately.

<b>Required Privilege</b>	routing—To view this statement in the configuration.
<b>Level</b>	routing-control—To add this statement to the configuration.

<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">container-label-switched-path on page 519</a></li></ul>
------------------------------	---

## traceoptions (Protocols RSVP)

<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> &lt;<i>flag-modifier</i>&gt; &lt;disable&gt;; } </pre>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable RSVP-level trace options.
<b>Default</b>	The default RSVP-level trace options are those inherited from the routing protocols <b>traceoptions</b> statement included at the [edit routing-options] hierarchy level.
<b>Options</b>	<p><b>disable</b>—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as <b>all</b>.</p> <p><b>filename</b>—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory <b>/var/log</b>. We recommend that you place RSVP tracing output in the file <b>rsvp-log</b>.</p> <p><b>files number</b>—(Optional) Maximum number of trace files. When a trace file named <b>trace-file</b> reaches its maximum size, it is renamed <b>trace-file.0</b>, then <b>trace-file.1</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>flag</b>—Tracing operation to perform. To specify more than one tracing operation, include multiple <b>flag</b> statements.</p> <ul style="list-style-type: none"> <li>• <b>all</b>—All tracing operations</li> <li>• <b>error</b>—All detected error conditions</li> <li>• <b>event</b>—RSVP-related events</li> <li>• <b>lmp</b>—RSVP-LMP interactions</li> <li>• <b>packets</b>—All RSVP packets</li> <li>• <b>path</b>—All path messages</li> <li>• <b>pathtear</b>—PathTear messages</li> </ul>

- **resv**—Resv messages
- **resvtear**—ResvTear messages
- **route**—Routing information
- **state**—Session state transitions, including when RSVP-signaled LSPs come up and go down.

**flag-modifier**—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

**no-world-readable**—(Optional) Enable only certain users to read the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** again reaches this size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of files.

**world-readable**—(Optional) Enable any user to read the log file.

<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>• <i>Tracing RSVP Protocol Traffic</i></li></ul>
------------------------------	--

## tunnel-services (RSVP)

<b>Syntax</b>	tunnel-services { devices <i>device-names</i> ; }
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp], [edit protocols rsvp]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Enable ultimate-hop popping on point-to-multipoint LSPs. The Junos OS selects one of the available virtual tunnel (VT) interfaces to de-encapsulate the egress traffic. By default, the selection process is performed automatically.
<b>Default</b>	Ultimate-hop popping is disabled.
<b>Options</b>	<b>devices</b> <i>device-names</i> —Specify which VT interfaces are used to handle the RSVP traffic. <b>Range:</b> 0 to 8 devices
<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>• <i>Enabling Ultimate-Hop Popping on Point-to-Multipoint LSPs</i></li> </ul>

## update-threshold

<b>Syntax</b>	update-threshold <i>threshold</i> ;
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> protocols rsvp interface <i>interface-name</i> ], [edit protocols rsvp interface <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
<b>Description</b>	Adjust the threshold at which a change in bandwidth triggers an interior gateway protocol (IGP) update.
<b>Options</b>	<b>threshold</b> —Specify the percentage change in bandwidth to trigger an IGP update. <b>Range:</b> 1 through 20 percent <b>Default:</b> 10 percent
<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>• <i>Configuring RSVP Interfaces</i></li> </ul>



## CHAPTER 9

# Monitoring Commands for RSVP

- `clear mpls container-lsp`
- `clear rsvp session`
- `clear rsvp statistics`
- `ping mpls rsvp`
- `request mpls container-lsp`
- `clear mpls container-lsp`
- `show rsvp interface`
- `show rsvp neighbor`
- `show rsvp session`
- `show rsvp statistics`
- `show rsvp version`
- `traceroute mpls rsvp`

## clear mpls container-lsp

---

<b>Syntax</b>	<code>clear mpls container-lsp</code> <code>&lt;autobandwidth&gt;</code> <code>&lt;history&gt;</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code> <code>&lt;member&gt;</code> <code>&lt;name <i>name</i>&gt;</code> <code>&lt;optimize   optimize-aggressive&gt;</code> <code>&lt;statistics&gt;</code>
<b>Release Information</b>	Statement introduced in Junos OS Release 14.2. Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.
<b>Description</b>	Release the routes and states associated with MPLS container label-switched paths (LSPs), and start new LSPs.
<b>Options</b>	<p><b>none</b>—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.</p> <p><b>autobandwidth</b>—(Optional) Clear LSP autobandwidth counters.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>name <i>name</i></b>—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the <i>Junos Network Interfaces Configuration Guide</i>.</p> <p><b>optimize   optimize-aggressive</b>—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.</p> <p><b>statistics</b>—(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (<b>name</b> and <b>path</b> options) on transit routers.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show mpls container-lsp</a></li><li>• <a href="#">request mpls container-lsp on page 562</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear mpls container-lsp on page 553</a> <a href="#">clear mpls container-lsp name on page 553</a> <a href="#">clear mpls container-lsp statistics on page 553</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.



## Sample Output

clear mpls container-lsp

```
user@host> clear mpls container-lsp
```

clear mpls container-lsp name

```
user@host> clear mpls container-lsp name name
```

clear mpls container-lsp statistics

```
user@host> clear mpls container-lsp statistics
```

## clear rsvp session

---

<b>List of Syntax</b>	<a href="#">Syntax on page 554</a> <a href="#">Syntax (EX and QFX Series Switches) on page 554</a>
<b>Syntax</b>	<pre>clear rsvp session &lt;connection-destination address&gt; &lt;connection-source address&gt; &lt;gracefully&gt; &lt;logical-system (all   logical-system-name)&gt; &lt;lsp-id identifier&gt; &lt;name name&gt; &lt;optimize-fast-reroute&gt; &lt;tunnel-id identifier&gt;</pre>
<b>Syntax (EX and QFX Series Switches)</b>	<pre>clear rsvp session &lt;connection-destination address&gt; &lt;connection-source address&gt; &lt;gracefully&gt; &lt;lsp-id identifier&gt; &lt;name name&gt; &lt;optimize-fast-reroute&gt; &lt;tunnel-id identifier&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
<b>Description</b>	Reset and restart Resource Reservation Protocol (RSVP) sessions.
<b>Options</b>	<p><b>none</b>—Reset and restart all RSVP sessions for which this routing device is the ingress, transit, or egress routing device.</p> <p><b>connection-source address</b>—(Optional) Source address for GMPLS and MPLS LSPs from the RSVP sender template.</p> <p><b>connection-destination address</b>—(Optional) Destination address for GMPLS and MPLS LSPs from the RSVP sender template.</p> <p><b>gracefully</b>—(Optional) Gracefully reset an RSVP session for a nonpacket LSP in two passes. In the first pass, the Admin-Status object is signaled along the path to the other endpoint of the RSVP session. In the second pass, the path used by the RSVP session is torn down. This option can only be used on the ingress or egress routing device of the RSVP session and is only valid for nonpacket LSPs.</p> <p><b>logical-system (all   logical-system-name)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>lsp-id identifier</b>—(Optional) LSP identifier (source port) for the RSVP sender template.</p> <p><b>name name</b>—(Optional) Reset and restart the specified RSVP session.</p> <p><b>optimize-fast-reroute</b>—(Optional) Begin fast reroute optimization.</p>

**tunnel-id *identifier***—(Optional) Tunnel identifier (destination port) for the RSVP session.

**Required Privilege Level**

clear

**Related Documentation**

- [clear mpls lsp on page 317](#)
- [show rsvp session on page 575](#)

**List of Sample Output** [clear rsvp session on page 555](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request.

## Sample Output

[clear rsvp session](#)

```
user@host> clear rsvp session
```

## clear rsvp statistics

---

<b>List of Syntax</b>	<a href="#">Syntax on page 556</a> <a href="#">Syntax (EX Series Switches) on page 556</a>
<b>Syntax</b>	clear rsvp statistics <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switches)</b>	clear rsvp statistics
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Clear Resource Reservation Protocol (RSVP) packet and error statistics.
<b>Options</b>	<b>none</b> —Clear RSVP packet and error statistics.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show rsvp statistics on page 585</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear rsvp statistics on page 556</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### clear rsvp statistics

```
user@host> clear rsvp statistics
```

## ping mpls rsvp

**Syntax** ping mpls rsvp  
 <lsp-name>  
 <count count>  
 <destination address>  
 <detail>  
 <dynamic-bypass>  
 <egress egress-address>  
 <exp forwarding-class>  
 <interface interface-name>  
 <logical-system (all | logical-system-name)>  
 <manual-bypass>  
 <multipoint>  
 <size bytes>  
 <source source-address>  
 <standby standby-path-name>  
 <sweep>

**Release Information** Command introduced before Junos OS Release 7.4.  
 The **egress** and **multipoint** options were introduced in Junos OS Release 9.2.  
 The **size** and **sweep** options were introduced in Junos OS Release 9.6.  
 The **dynamic-bypass** and **manual-bypass** options were introduced in Junos OS Release 10.2.  
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

**Description** Check the operability of MPLS RSVP-signaled label-switched path (LSP) connections. Type Ctrl+c to interrupt a **ping mpls** command.

**Options** **count count**—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

**destination address**—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

**detail**—(Optional) Display detailed information about the echo requests sent and received.



**NOTE:** When using the **detail** option, the reported time is based on the system time configured on the local and remote routers. Differences in these system times can result in inaccurate one way ping trip times being reported.

In practice, it is difficult to synchronize the system times of independent Juniper Networks routers with sufficient accuracy to provide a meaningful time value for the **detail** option (even when synchronized using NTP).

**dynamic-bypass**—(Optional) Ping dynamically generated bypass LSPs, used for protecting other LSPs.

**egress *egress-address***—(Optional) Only the specified egress router or switch responds to the ping request.

**exp *forwarding-class***—(Optional) Value of the forwarding class for the MPLS ping packets.

**interface**—(Optional) Specify the name of the interface protected by the manual bypass LSP. This option is only available when you have also used the **manual-bypass** option.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on the specified logical system.

***lsp-name***—Ping an RSVP-signaled LSP using an LSP name.

**manual-bypass**—(Optional) Ping manually configured bypass LSPs, used for protecting other LSPs. For this option, you must also specify the interface protected by the manual bypass LSP using the **interface** option.

**multipoint**—(Optional) Send ping requests to each of the egress routers or switches participating in a point-to-multipoint LSP. You can also include the **egress** option to ping a specific egress router or switch participating in a point-to-multipoint LSP.

**size *bytes***—(Optional) Size of the LSP ping request packet (100 through 65468 bytes). Packets are 4-byte aligned. For example, if you enter a size of 101, 102, 103, or 104, the router or switch uses a size value of 104 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 100-byte minimum.

**source *source-address***—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface.

**standby *standby-path-name***—(Optional) Name of the standby path.

**sweep**—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

**Additional Information** If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

**Required Privilege Level** network

**List of Sample Output** [ping mpls rsvp \(Echo Reply Received\) on page 559](#)  
[ping mpls rsvp \(Echo Reply with Error Code\) on page 559](#)

[ping mpls rsvp detail on page 559](#)

[ping mpls rsvp multipoint egress detail count on page 559](#)

[ping mpls rsvp multipoint detail count on page 559](#)

[ping mpls rsvp destination detail count size on page 560](#)

[ping mpls rsvp destination detail sweep size on page 560](#)

**Output Fields** When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

## Sample Output

### ping mpls rsvp (Echo Reply Received)

```
user@host> ping mpls rsvp test1
!!!!!--- lsping statistics ---5 packets transmitted, 5 packets received, 0% packet
loss
```

### ping mpls rsvp (Echo Reply with Error Code)

```
user@host> ping mpls rsvp test2
!!xxx--- lsping statistics ---5 packets transmitted, 2 packets received, 60%
packet loss3 packets received with error status, not counted as received.
```

### ping mpls rsvp detail

```
user@host> ping mpls rsvp to-green detail
Request for seq 1, to interface 67, labels <100095, 0, 0>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, labels <100095, 0, 0>
Reply for seq 2, return code: Egress-ok
```

### ping mpls rsvp multipoint egress detail count

```
user@host>ping mpls rsvp sample-lsp multipoint egress 192.168.1.3 detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.242 ms
Local transmit time: 1205310695s 215737us
Remote receive time: 1205310695s 215979us

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

### ping mpls rsvp multipoint detail count

```
user@host>ping mpls rsvp sample-lsp multipoint detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, return code: Unknown TLV, time: 9.877 m Local transmit time:
1205310615s 347317us
Remote receive time: 1205310615s 357194us
Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.351 ms
```

```

Local transmit time: 1205310615s 347262us
Remote receive time: 1205310615s 347613us
Reply for seq 1, egress 192.168.1.13, return code: Egress-ok, time: 0.301 ms
Local transmit time: 1205310615s 347167us
Remote receive time: 1205310615s 347468us
Timeout for seq 1, egress 192.168.1.1
Timeout for seq 1, egress 192.168.1.4
Timeout for seq 1, egress 192.168.1.14

--- lsping, egress 192.168.1.1 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.4 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.13 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.14 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

```

#### ping mpls rsvp destination detail count size

```

user@host> ping mpls rsvp chaser-access destination 192.168.0.1 detail count 1 size 4468

Request for seq 1, to interface 88, label 299984, packet size 4468
Reply for seq 1, return code: Egress-ok, time: 44.804 ms
    Local transmit time: 2009-03-30 22:05:02 CEST 408.629 ms
    Remote receive time: 2009-03-30 22:05:02 CEST 453.433 ms

--- lsping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

```

#### ping mpls rsvp destination detail sweep size

```

user@router> ping mpls rsvp chaser-access destination 192.168.0.1 detail sweep size 4500
Request for seq 1, to interface 86, no label stack., packet size 100
Reply for seq 1, return code: Egress-ok, time: -39.264 ms
    Local transmit time: 2009-04-24 14:05:40 CEST 541.423 ms
    Remote receive time: 2009-04-24 14:05:40 CEST 502.159 ms
Request for seq 2, to interface 86, no label stack., packet size 2300
Reply for seq 2, return code: Egress-ok, time: -38.179 ms
    Local transmit time: 2009-04-24 14:05:41 CEST 544.240 ms
    Remote receive time: 2009-04-24 14:05:41 CEST 506.061 ms
Request for seq 3, to interface 86, no label stack., packet size 4500
Timeout for seq 3
Request for seq 4, to interface 86, no label stack., packet size 3400
Reply for seq 4, return code: Egress-ok, time: -37.545 ms
    Local transmit time: 2009-04-24 14:05:45 CEST 549.953 ms
    Remote receive time: 2009-04-24 14:05:45 CEST 512.408 ms
Request for seq 5, to interface 86, no label stack., packet size 3952
Reply for seq 5, return code: Egress-ok, time: -37.176 ms
    Local transmit time: 2009-04-24 14:05:46 CEST 555.881 ms
    Remote receive time: 2009-04-24 14:05:46 CEST 518.705 ms
Request for seq 6, to interface 86, no label stack., packet size 4228
Reply for seq 6, return code: Egress-ok, time: -36.962 ms
    Local transmit time: 2009-04-24 14:05:47 CEST 561.809 ms
    Remote receive time: 2009-04-24 14:05:47 CEST 524.847 ms

```



```
Request for seq 7, to interface 86, no label stack., packet size 4368
Reply for seq 7, return code: Egress-ok, time: -36.922 ms
    Local transmit time: 2009-04-24 14:05:48 CEST 568.738 ms
    Remote receive time: 2009-04-24 14:05:48 CEST 531.816 ms
Request for seq 8, to interface 86, no label stack., packet size 4440
Reply for seq 8, return code: Egress-ok, time: -36.855 ms
    Local transmit time: 2009-04-24 14:05:49 CEST 575.669 ms
    Remote receive time: 2009-04-24 14:05:49 CEST 538.814 ms
Request for seq 9, to interface 86, no label stack., packet size 4476
Timeout for seq 9
Request for seq 10, to interface 86, no label stack., packet size 4460
Reply for seq 10, return code: Egress-ok, time: -36.906 ms
    Local transmit time: 2009-04-24 14:05:53 CEST 584.382 ms
    Remote receive time: 2009-04-24 14:05:53 CEST 547.476 ms
Request for seq 11, to interface 86, no label stack., packet size 4480
Timeout for seq 11
Request for seq 12, to interface 86, no label stack., packet size 4472
Timeout for seq 12
Request for seq 13, to interface 86, no label stack., packet size 4468
Reply for seq 13, return code: Egress-ok, time: -36.943 ms
    Local transmit time: 2009-04-24 14:06:00 CEST 594.884 ms
    Remote receive time: 2009-04-24 14:06:00 CEST 557.941 ms
Request for seq 14, to interface 86, no label stack., packet size 4476
Timeout for seq 14
Request for seq 15, to interface 86, no label stack., packet size 4472
Timeout for seq 15

--- lsp ping sweep result---
Maximum Transmission Unit (MTU) is 4468 bytes
```

## request mpls container-lsp

---

<b>Syntax</b>	<code>request mpls container-lsp</code> <code>&lt;logical-system (all   <i>logical-system-name</i>)&gt;</code> <code>&lt;name <i>lsp-name</i>&gt;</code> <code>&lt;adjust-autobandwidth&gt;</code> <code>&lt;normalization&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 14.2. Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.
<b>Description</b>	Manually trigger a bandwidth allocation adjustment for the container label-switched path (LSP).
<b>Options</b>	<p><b>none</b>—Manually trigger a bandwidth allocation adjustment for all active member LSP paths.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>name <i>lsp-name</i></b>—(Optional) Manually trigger a bandwidth allocation adjustment on the specified member LSP only.</p> <p><b>adjust-autobandwidth</b>—(Optional) Request LSP autobandwidth adjustment.</p> <p><b>normalization</b>—(Optional) Request container LSP normalization.</p>
<b>Required Privilege Level</b>	clear, maintenance
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show mpls container-lsp</a></li><li>• <a href="#">clear mpls container-lsp on page 552</a></li></ul>
<b>List of Sample Output</b>	<a href="#">request mpls container-lsp on page 562</a> <a href="#">request mpls container-lsp on page 562</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request mpls container-lsp

```
user@host> request mpls container-lsp lsp-name normalize
```

### request mpls container-lsp

```
user@host> request mpls container-lsp normalize bandwidth bps
```

## clear mpls container-lsp

<b>Syntax</b>	<pre>clear mpls container-lsp &lt;autobandwidth&gt; &lt;history&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;member&gt; &lt;name <i>name</i>&gt; &lt;optimize   optimize-aggressive&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 14.2.</p> <p>Statement introduced for QFX Switches in Junos OS Release 15.1X53-D30.</p>
<b>Description</b>	Release the routes and states associated with MPLS container label-switched paths (LSPs), and start new LSPs.
<b>Options</b>	<p><b>none</b>—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.</p> <p><b>autobandwidth</b>—(Optional) Clear LSP autobandwidth counters.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><b>name <i>name</i></b>—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the <i>Junos Network Interfaces Configuration Guide</i>.</p> <p><b>optimize   optimize-aggressive</b>—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.</p> <p><b>statistics</b>—(Optional) Clear LSP statistics. You cannot clear the MPLS LSP statistics using a regular expression (<b>name</b> and <b>path</b> options) on transit routers.</p>
<b>Required Privilege Level</b>	clear
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show mpls container-lsp</a></li> <li>• <a href="#">request mpls container-lsp on page 562</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">clear mpls container-lsp on page 564</a></p> <p><a href="#">clear mpls container-lsp name on page 564</a></p> <p><a href="#">clear mpls container-lsp statistics on page 564</a></p>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

clear mpls container-lsp

```
user@host> clear mpls container-lsp
```

clear mpls container-lsp name

```
user@host> clear mpls container-lsp name name
```

clear mpls container-lsp statistics

```
user@host> clear mpls container-lsp statistics
```

## show rsvp interface

<b>List of Syntax</b>	<a href="#">Syntax on page 565</a> <a href="#">Syntax (EX Series Switches) on page 565</a>
<b>Syntax</b>	<pre>show rsvp interface &lt;brief   detail   extensive&gt; &lt;instance <i>instance-name</i>&gt; &lt;link-management&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show rsvp interface &lt;brief   detail   extensive&gt; &lt;link-management&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display the status of Resource Reservation Protocol (RSVP)-enabled interfaces and packet statistics.
<b>Options</b>	<p><b>none</b>—Display standard information about the status of RSVP-enabled interfaces and packet statistics.</p> <p><b>brief   detail   extensive   link-management</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display RSVP status information for the specified instance. If <i>instance-name</i> is omitted, RSVP status information is displayed for the master instance.</p> <p><b>link-management</b>—(Optional) Use the link-management option to display the control peers and corresponding TE-link information created by the Link Management Protocol (LMP).</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show rsvp interface brief on page 568</a> <a href="#">show rsvp interface detail on page 568</a> <a href="#">show rsvp interface extensive on page 568</a> <a href="#">show rsvp interface link-management on page 569</a>
<b>Output Fields</b>	Table 65 on page 566 lists the output fields for the <b>show rsvp interface</b> command. Output fields are listed in the approximate order in which they appear.

Table 65: show rsvp interface Output Fields

Field Name	Field Description	Level of Output
<b>RSVP interface</b>	Number of interfaces on which RSVP is active. Each interface has one line of output.	All levels
<b>Interface</b>	Name of the interface.	All levels
<b>Index</b>	Index of the interface.	<b>detail</b>
<b>State</b>	State of the interface. <ul style="list-style-type: none"> <li>• <b>Disabled</b>—No traffic engineering information is displayed.</li> <li>• <b>Down</b>—Interface is not operational.</li> <li>• <b>Enabled</b>—Displays traffic engineering information.</li> <li>• <b>Up</b>—Interface is operational.</li> </ul>	All levels
<b>NoAuthentication</b>	Interface does not support RSVP authentication.	<b>detail</b>
<b>NoAggregate</b>	Interface does not support refresh reduction.	<b>detail</b>
<b>NoReliable</b>	Interface does not support refresh reduction message ID extension.	<b>detail</b>
<b>NoLinkProtection</b>	Interface does not support link protection.	<b>detail</b>
<b>HelloInterval</b>	Frequency at which RSVP hellos are sent on this interface (in seconds).	<b>detail</b>
<b>Address</b>	IP address of the local interface.	<b>detail</b>
<b>Active control channel</b>	Next-hop link address to transmit messages.	None specified
<b>TELink</b>	Traffic-engineered links that are managed by the peer they are associated with.	None specified
<b>Active resv</b>	Number of reservations that are actively reserving bandwidth on the interface.	All levels
<b>PreemptionCnt</b>	Number of times an RSVP session was preempted on this interface.	<b>detail</b>
<b>Update threshold</b>	Percentage change in reserved bandwidth to trigger an IGP update.	<b>detail</b>
<b>Subscription</b>	User-configured subscription factor.	All levels
<b>bc number</b>	Bandwidth allocated for the specified bandwidth constraint.	<b>extensive</b>
<b>ct number</b>	Bandwidth allocated for the specified class type.	<b>extensive</b>
<b>Static BW</b>	Total interface bandwidth, in bps.	All levels
<b>Available BW</b>	Amount of bandwidth that RSVP is allowed to reserve, in bps. It is equal to (static bandwidth * subscription factor).	al levels

Table 65: show rsvp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Reserved BW</b>	Currently reserved bandwidth, in bps.	All levels
<b>SoftPreemptionCnt</b>	Number of times a soft preemption occurred on this interface. This number is not included in the <b>PreemptionCnt</b> value.	<b>detail</b>
<b>Overbooked BW</b>	Currently overbooked bandwidth, in bps, by class type ( <b>ct0</b> through <b>ct3</b> ).	<b>detail</b>
<b>Highwater mark</b>	Highest bandwidth that has ever been reserved on this interface, in bps.	<b>brief</b>
<b>PacketType</b>	Type of RSVP packet.	<b>detail</b>
<b>Total Sent</b>	Total number of packets sent.	<b>detail</b>
<b>Total Received</b>	Total number of packets received since RSVP was enabled.	<b>detail</b>
<b>Last 5 seconds Sent</b>	Number of packets sent in the last 5 seconds.	<b>detail</b>
<b>Last 5 seconds Received</b>	Number of packets received in the last 5 seconds.	<b>detail</b>
<b>Path</b>	Statistics about Path messages, which are sent from the RSVP sender along the data paths and store path state information in each node along the path.	<b>detail</b>
<b>PathErr</b>	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.	<b>detail</b>
<b>PathTear</b>	Statistics about PathTear messages, which remove path states and dependent reservation states in any routers along a path.	<b>detail</b>
<b>Resv</b>	Statistics about Resv messages, which are sent from the RSVP receiver along the data paths and store reservation state information in each node along the path.	<b>detail</b>
<b>ResvErr</b>	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.	<b>detail</b>
<b>ResvTear</b>	Statistics about ResvTear messages, which remove reservation states along a path.	<b>detail</b>
<b>Hello</b>	Number of RSVP hello packets that have been sent to and received from the neighbor.	<b>detail</b>
<b>Ack</b>	Acknowledge message for refresh reductions.	<b>detail</b>
<b>Srefresh</b>	Summary refresh messages.	<b>detail</b>
<b>EndtoEnd RSVP</b>	Statistics for the number of end-to-end RSVP messages sent.	<b>detail</b>
<b>Queue</b>	CoS transmit queue number and its associated forwarding class designation.	<b>extensive</b>

Table 65: show rsvp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>TxRate</b>	Configured bandwidth in Mbps and configured bandwidth as a percentage of the specified queue.	<b>extensive</b>
<b>Priority</b>	Weight of the queue relative to other configured queues, in percentage.	<b>extensive</b>
<i>queue-priority-value</i>	<b>Low, High, None, or Exact.</b> <b>None</b> indicates no rate limiting. <b>Exact</b> indicates the queue transmits at the configured rate only.	<b>extensive</b>

## Sample Output

### show rsvp interface brief

```

user@host> show rsvp interface brief
RSVP interface: 1 active

```

Interface	State	Active resv	Subscr- ption	Static BW	Available BW	Reserved BW	Highwater mark
de0.0	Up	1	23%	10Mbps	989.992kbps	1.31Mbps	1.31Mbps

### show rsvp interface detail

```

user@host> show rsvp interface detail
so-0/1/1.0 Index 6, State: Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 3(second)
  Address 192.168.207.29, 10.255.245.194
  ActiveResv 0, PreemptionCnt 0, Update threshold 10%
  Subscription 100%, StaticBW 155.52Mbps, AvailableBW 155.52Mbps
  ReservedBW [0] 155Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  SoftPreemptionCnt1
  OverbookedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 155Mbps[5] 0bps[6] 0bps[7] 0bps
  PacketType
    Total
    Last 5 seconds
    Sent Received Sent Received
    Path 16 0 1 0
    PathErr 0 0 0 0
    PathTear 1 0 0 0
    Resv 0 11 0 1
    ResvErr 0 0 0 0
    ResvTear 0 0 0 0
    Hello 66 67 1 1
    Ack 0 0 0 0
    Srefresh 0 0 0 0
    EndtoEnd RSVP 0 0 0 0
    ...

```

### show rsvp interface extensive

```

user@host> show rsvp interface extensive
so-1/0/0.0 Index 72, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 192.168.213.22, 10.255.240.175
  ActiveResv 1, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = (ct0+ct1+ct2+ct3), StaticBW 622.08Mbps
  bc1 = (ct1+ct2+ct3), StaticBW 466.56Mbps

```



```

bc2 = (ct2+ct3), StaticBW 311.04Mbps
bc3 = ct3, StaticBW 155.52Mbps
ct0: StaticBW 155.52Mbps, AvailableBW 522.08Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
ct1: StaticBW 155.52Mbps, AvailableBW 366.56Mbps
ReservedBW [0] 100Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps

ct2: StaticBW 155.52Mbps, AvailableBW 311.04Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
ct3: StaticBW 155.52Mbps, AvailableBW 155.52Mbps
ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
Queue          TxRate          Priority Exact
0              155.52Mbps          25%      Low
1              155.52Mbps          25%      Low
2              155.52Mbps          25%      Low
3              155.52Mbps          25%      Low

```

### show rsvp interface link-management

```

user@host> show rsvp interface link-management
RSVP interface: 2 active
PEER-C State: Up
Active Control Channel: so-0/1/0.0

TElink: TElnk1, Link ID: 37811
ActiveResv 0, PreemptionCnt 0
StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

TElink: TElnk2, Link ID: 37808
ActiveResv 1, PreemptionCnt 0
StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

PEER-B State: Up
Active Control Channel: so-1/0/0.0

TElink: TElnkAB1, Link ID: 1598
ActiveResv 0, PreemptionCnt 0
StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps

TElink: TElnkAB2, Link ID: 1597
ActiveResv 0, PreemptionCnt 0
StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps

```

## show rsvp neighbor

<b>List of Syntax</b>	<a href="#">Syntax on page 570</a> <a href="#">Syntax (EX Series Switches) on page 570</a>
<b>Syntax</b>	<pre>show rsvp neighbor &lt;brief   detail&gt; &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	<pre>show rsvp neighbor &lt;brief   detail&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display Resource Reservation Protocol (RSVP) neighbors that were discovered dynamically during the exchange of RSVP packets.
<b>Options</b>	<p><b>none</b>—Display standard information about RSVP neighbors.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display the RSVP neighbor information for the specified instance. If <i>instance-name</i> is omitted, RSVP neighbor information is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show rsvp neighbor on page 574</a> <a href="#">show rsvp neighbor detail on page 574</a>
<b>Output Fields</b>	<a href="#">Table 66 on page 570</a> lists the output fields for the <b>show rsvp neighbor</b> command. Output fields are listed in the approximate order in which they appear.

**Table 66: show rsvp neighbor Output Fields**

Field Name	Field Description	Level of Output
<b>RSVP neighbor</b>	Number of neighbors that the routing device has learned of. Each neighbor has one line of output.	All levels
<b>via</b>	Name of the interface where the neighbor has been detected. In the case of generalized MPLS (GMPLS) LSPs, the name of the peer where the neighbor has been detected.	<b>detail</b>
<b>Address</b>	Address of a learned neighbor.	All levels

Table 66: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Idle</b>	Length of time the neighbor has been idle, in seconds.	All levels
<b>Up/Dn</b>	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.	All levels
<b>Up cnt and Down cnt</b>	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets, such as routers running Junos OS Release 3.2 or earlier, are not reported as up or down.	detail
<b>status</b>	State of the RSVP neighbor: <ul style="list-style-type: none"> <li>• <b>Up</b>—Routing device can detect RSVP Hello messages from the neighbor.</li> <li>• <b>Down</b>—Routing device has received one of the following indications:               <ul style="list-style-type: none"> <li>• Communication failure from the neighbor.</li> <li>• Communication from IGP that the neighbor is unavailable.</li> <li>• Change in the sequence numbers in the RSVP Hello messages sent by the neighbor.</li> </ul> </li> <li>• <b>Restarting</b>—RSVP neighbor is unavailable and might be restarting. The neighbor remains in this state until it has restarted or is declared dead. This state is possible only when graceful restart is enabled.</li> <li>• <b>Restarted</b>—RSVP neighbor has restarted and is undergoing state recovery (graceful restart) procedures.</li> <li>• <b>Dead</b>—Routing device has lost all communication with the RSVP neighbor. Any RSVP sessions with that neighbor are torn down.</li> </ul>	detail
<b>LastChange</b>	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is <b>hh:mm:ss</b> .	All levels
<b>Last changed time</b>	Time elapsed since the neighbor state changed either from up to down or from down to up.	detail
<b>HelloInt</b>	Frequency at which RSVP hellos are sent on this interface (in seconds).	All levels
<b>HelloTx/Rx</b>	Number of hello packets sent to and received from the neighbor.	All levels
<b>Hello</b>	Number of RSVP hello packets that have been sent to and received from the neighbor.	detail
<b>Message received</b>	Number of Path and Resv messages that this routing device has received from the neighbor.	detail
<b>Remote Instance</b>	Identification provided by the remote routing device during Hello message exchange.	detail

Table 66: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Local Instance</b>	Identification sent to the remote routing device during Hello message exchange.	<b>detail</b>
<b>Refresh reduction</b>	<p>Measure of processing overhead requests of refresh messages. Refresh reduction extensions improve routing device performance by reducing the process overhead, thus increasing the number of LSPs a routing device can support. <b>Refresh reduction</b> can have the following values:</p> <ul style="list-style-type: none"> <li>• <b>operational</b>—All four RSVP refresh reduction extensions—message ack, bundling, summary refresh, and staged refresh timer—are functional between the two neighboring routing devices. For a detailed explanation of these extensions, see RFC 2961.</li> <li>• <b>incomplete</b>—Some RSVP refresh reduction extensions are functional between the two neighboring routing devices.</li> <li>• <b>no operational</b>—Either the refresh reduction feature has been turned off, or the remote routing device cannot support the refresh reduction extensions.</li> </ul>	<b>detail</b>
<b>Remote end</b>	<p>Neighboring routing device's status with regard to refresh reduction:</p> <ul style="list-style-type: none"> <li>• <b>enabled</b>—Remote routing device has requested refresh reduction during RSVP message exchanges.</li> <li>• <b>disabled</b>—Remote routing device does not require refresh reduction.</li> </ul>	<b>detail</b>
<b>Ack-extension</b>	<p>An RSVP refresh reduction extension:</p> <ul style="list-style-type: none"> <li>• <b>enabled</b>—Both local and remote routing devices support the ack-extension (RFC 2961).</li> <li>• <b>disabled</b>—Remote routing device does not support the ack-extension.</li> </ul>	<b>detail</b>
<b>Link protection</b>	<p>Status of the MPLS fast reroute mechanism that protects traffic from link failure:</p> <ul style="list-style-type: none"> <li>• <b>enabled</b>—Link protection feature has been turned on, protecting the neighbor with a bypass LSP.</li> <li>• <b>disabled</b>—No link protection feature has been enabled for this neighbor.</li> </ul>	<b>detail</b>
<b>LSP name</b>	Name of the bypass LSP.	<b>detail</b>
<b>Bypass LSP</b>	<p>Status of the bypass LSP. It can have the following values:</p> <ul style="list-style-type: none"> <li>• <b>does not exist</b>—Bypass LSP is not available.</li> <li>• <b>connecting</b>—Routing device is in the process of establishing a bypass LSP, and the LSP is not available for link protection at the moment.</li> <li>• <b>operational</b>—Bypass LSP is up and running.</li> <li>• <b>down</b>—Bypass LSP has gone down, with the most probable cause a node or a link failure on the bypass path.</li> </ul>	<b>detail</b>
<b>Backup routes</b>	Number of user LSPs (or routes) that are being protected by a bypass LSP (before link failure).	<b>detail</b>
<b>Backup LSPs</b>	Number of LSPs that have been temporarily established to maintain traffic by refreshing the downstream LSPs during link failure (not a one-to-one correspondence).	<b>detail</b>

Table 66: show rsvp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Bypass explicit route</b>	Explicit route object's (ERO) path that is taken by the bypass LSP.	<b>detail</b>
<b>Restart time</b>	Length of time a neighbor waits to receive a Hello from the restarting node before declaring the node dead and deleting the states (in milliseconds).	<b>detail</b>
<b>Recovery time</b>	Length of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds). Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.	<b>detail</b>

## Sample Output

### show rsvp neighbor

```
user@host> show rsvp neighbor
RSVP neighbor: 2 learned
Address          Idle Up/Dn  LastChange HelloInt HelloTx/Rx
192.168.207.203   0  3/2    13:01      3   366/349
192.168.207.207   0  1/0    22:49      3   448/448
```

### show rsvp neighbor detail

```
user@host> show rsvp neighbor detail
RSVP neighbor: 2 learned
Address: 192.168.207.203   via: ecstasy1 status: Up
  Last changed time: 28:47, Idle: 0 sec, Up cnt: 3, Down cnt: 2
  Message received: 632
  Hello: sent 673, received 656, interval 3 sec
  Remote instance: 0x6432838a, Local instance: 0x74b72e36
  Refresh reduction: operational
    Remote end: enabled, Ack-extension: enabled
  Link protection: enabled
    LSP name: Bypass_to_192.168.207.203
    Bypass LSP: operational, Backup routes: 1, Backup LSPs: 0
    Bypass explicit route: 192.168.207.207 192.168.207.224
  Restart time: 60000 msec, Recovery time: 0 msec
```

## show rsvp session

<b>List of Syntax</b>	<a href="#">Syntax on page 575</a> <a href="#">Syntax (EX and QFX Series Switches) on page 575</a>
<b>Syntax</b>	<pre>show rsvp session &lt;brief   detail   extensive   terse&gt; &lt;bidirectional   unidirectional&gt; &lt;bypass&gt; &lt;down   up&gt; &lt;externally-provisioned&gt; &lt;instance <i>instance-name</i>&gt; &lt;interface <i>interface-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt; &lt;lsp-type&gt; &lt;name <i>session-name</i>&gt; &lt;p2mp&gt; &lt;session-type&gt; &lt;statistics&gt; &lt;te-link <i>te-link</i>&gt;</pre>
<b>Syntax (EX and QFX Series Switches)</b>	<pre>show rsvp session &lt;brief   detail   extensive   terse&gt; &lt;bidirectional   unidirectional&gt; &lt;bypass&gt; &lt;down   up&gt; &lt;externally-provisioned&gt; &lt;interface <i>interface-name</i>&gt; &lt;lsp-type&gt; &lt;name <i>session-name</i>&gt; &lt;p2mp&gt; &lt;session-type&gt; &lt;statistics&gt; &lt;te-link <i>te-link</i>&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>externally-provisioned</b> option added in Junos OS Release 13.3.</p> <p>Command introduced in Junos OS Release 13.2X51-D15 for QFX Series.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display information about Resource Reservation Protocol (RSVP) sessions.
<b>Options</b>	<p><b>none</b>—Display standard information about all RSVP sessions.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>bidirectional   unidirectional</b>—(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.</p> <p><b>bypass</b>—(Optional) Display RSVP sessions for bypass LSPs.</p> <p><b>down   up</b>—(Optional) Display only LSPs that are inactive or active, respectively.</p>

**externally-provisioned**—(Optional) Display the LSPs that are generated dynamically and provisioned by an external Path Computation Element (PCE).

**instance *instance-name***—(Optional) Display RSVP sessions for the specified instance. If *instance-name* is omitted, RSVP session information is displayed for the master instance.

**interface *interface-name***—(Optional) Display RSVP sessions for the specified interface only.

**logical-system (all | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.

***lsp-type***—(Optional) Display information about RSVP sessions with regard to LSPs:

- **bypass**—Sessions used for bypass LSPs.
- **lsp**—Sessions used to set up LSPs.
- **nolsp**—Sessions not used to set up LSPs.

**name *session-name***—(Optional) Display information about the named session.

**p2mp**—(Optional) Display point-to-multipoint information.

***session-type***—(Optional) Display information about a particular session type:

- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that transit through this routing device.

**statistics**—(Optional) Display packet statistics.

**te-link *te-link***—(Optional) Display sessions with reservations on the specified TE link.

**Required Privilege  
Level**

view

**Related  
Documentation**

- [clear rsvp session on page 554](#)

**List of Sample Output**

[show rsvp session on page 580](#)  
[show rsvp session statistics on page 580](#)  
[show rsvp session detail on page 581](#)  
[show rsvp session detail \(When Egress Protection is in Standby Mode\) on page 581](#)  
[show rsvp session detail \(When Egress Protection is in Effect During a Local Repair\) on page 581](#)  
[show rsvp session detail \(Path MTU Output Field\) on page 582](#)  
[show rsvp session detail \(GMPLS\) on page 582](#)  
[show rsvp session extensive on page 582](#)  
[show rsvp session p2mp \(Ingress Router\) on page 583](#)  
[show rsvp session p2mp \(Transit Router\) on page 583](#)



**Output Fields** Table 67 on page 577 describes the output fields for the **show rsvp session** command. Output fields are listed in the approximate order in which they appear.

**Table 67: show rsvp session Output Fields**

Field Name	Field Description	Level of Output
<b>Ingress RSVP</b>	Information about ingress RSVP sessions.	<b>detail</b>
<b>Ingress RSVP</b>	Information about ingress RSVP sessions. Each session has one line of output.	All levels
<b>Egress RSVP</b>	Information about egress RSVP sessions.	All levels
<b>Transit RSVP</b>	Information about the transit RSVP sessions.	All levels
<b>P2MP name</b>	(Appears only when the <b>p2mp</b> option is specified). Name of the point-to-multipoint LSP path.	All levels
<b>P2MP branch count</b>	(Appears only when the <b>p2mp</b> option is specified). Number of LSPs receiving packets from the point-to-multipoint LSP.	All levels
<b>To</b>	Destination (egress routing device) of the session.	All levels
<b>From</b>	Source (ingress routing device) of the session.	All levels
<b>State</b>	State of the path: <b>Up</b> , <b>Down</b> , or <b>AdminDn</b> . <b>AdminDn</b> indicates that the LSP is being taken down gracefully.	All levels
<b>Address</b>	Destination (egress routing device) of the LSP.	<b>detail</b>
<b>From</b>	Source (ingress routing device) of the session.	<b>detail</b>
<b>LSPstate</b>	State of the LSP that is being handled by this RSVP session. It can be either <b>Up</b> , <b>Dn</b> (down), or <b>AdminDn</b> . <b>AdminDn</b> indicates that the LSP is being taken down gracefully.	<b>brief detail</b>
<b>Rt</b>	Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table ( <b>inet.0</b> ). For transit and egress RSVP sessions, the routing table is the primary MPLS table ( <b>mpls.0</b> ).	<b>brief</b>
<b>Active Route</b>	Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table ( <b>inet.0</b> ). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table ( <b>mpls.0</b> ).	<b>detail</b>
<b>LSPname</b>	Name of the LSP.	<b>brief detail</b>
<b>LSPpath</b>	Indicates whether the RSVP session is for the primary or secondary LSP path. <b>LSPpath</b> can be either <b>primary</b> or <b>secondary</b> and can be displayed on the ingress, egress, and transit routing devices. <b>LSPpath</b> can also indicate when a graceful LSP deletion has been triggered.	<b>detail</b>

Table 67: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Bypass</b>	(Egress routing device) Destination address for the bypass LSP.	<b>detail</b>
<b>Bidir</b>	(When LSP is bidirectional) LSP will allow data to travel in both directions between GMPLS devices.	<b>detail</b>
<b>Bidirectional</b>	(When LSP is bidirectional) LSP will allow data to travel both ways between GMPLS devices.	<b>detail</b>
<b>Upstream label in</b>	(When LSP is bidirectional) Incoming label for reverse direction traffic for this LSP.	<b>detail</b>
<b>Upstream label out</b>	(When LSP is bidirectional) Outgoing label for reverse direction traffic for this LSP.	<b>detail</b>
<b>Recovery label received</b>	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	<b>detail</b>
<b>Recovery label sent</b>	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv messages that is returned.	<b>detail</b>
<b>Suggested label received</b>	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	<b>detail</b>
<b>Suggested label sent</b>	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv message that is returned.	<b>detail</b>
<b>Resv style or Style</b>	RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be <b>FF</b> (fixed filter), <b>SE</b> (shared explicit), or <b>WF</b> (wildcard filter).	<b>brief detail</b>
<b>Label in</b>	Incoming label for this LSP.	<b>brief detail</b>
<b>Label out</b>	Outgoing label for this LSP.	<b>brief detail</b>
<b>Time left</b>	Number of seconds remaining in the lifetime of the reservation.	<b>brief detail</b>
<b>Since</b>	Date and time when the RSVP session was initiated.	<b>detail</b>
<b>Tspec</b>	Sender's traffic specification, which describes the sender's traffic parameters.	<b>detail</b>
<b>DiffServ info</b>	Indicates whether the LSP is a multiclass LSP ( <b>multiclass diffServ-TE LSP</b> ) or a Differentiated-Services-aware traffic engineering LSP ( <b>diffServ-TE LSP</b> ).	<b>detail</b>
<b>bandwidth</b>	Bandwidth for each class type ( <b>ct0</b> , <b>ct1</b> , <b>ct2</b> , or <b>ct3</b> ).	<b>detail</b>
<b>Port number</b>	Protocol ID and sender/receiver port used in this RSVP session.	<b>detail</b>
<b>Attrib flags</b>	<b>Non-PHP</b> indicates that ultimate hop popping has been requested by the LSP using this RSVP session	<b>extensive</b>

Table 67: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>FastReroute desired</b>	Fast reroute has been requested by the ingress routing device.	<b>detail</b>
<b>Soft preemption desired</b>	Soft preemption has been requested by the ingress routing device.	<b>detail</b>
<b>FastReroute desired</b>	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Fast reroute (one-to-one backup) has been requested by the ingress routing device.	<b>detail extensive</b>
<b>Link protection desired</b>	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Link protection (many-to-one backup) has been requested by the ingress routing device.	<b>detail extensive</b>
<b>Node/Link protection desired</b>	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Node and link protection (many-to-one backup) has been requested by the ingress routing device.	<b>detail extensive</b>
<b>Type</b>	<p>LSP type:</p> <ul style="list-style-type: none"> <li>• <b>Link protected LSP</b>—LSP has been protected by link protection at the outgoing interface. The name of the bypass used is also listed here (<b>extensive</b>).</li> <li>• <b>Node/Link protected LSP</b>—LSP has been protected by node and link protection at the outgoing interface. The name of the bypass used is also listed here (<b>extensive</b>).</li> <li>• <b>Protection down</b>—LSP is not currently protected.</li> <li>• <b>Bypass LSP</b>—LSP that is used to protect one or more user LSPs in case of link failure.</li> <li>• <b>Backup LSP at Point-of-Local-Repair (PLR)</b>—LSP that has been temporarily established to protect a user LSP at the ingress of a failed link.</li> <li>• <b>Backup LSP at Merge Point (MP)</b>—LSP that has been temporarily established to protect a user LSP at the egress of a failed link.</li> </ul>	<b>detail extensive</b>
<b>New bypass</b>	New bypass (the bypass name is also displayed) has been activated to protect the LSP.	<b>extensive</b>
<b>Link protection up, using <i>bypass-name</i></b>	Link protection (the bypass name is also displayed) has been activated for the LSP.	<b>extensive</b>
<b>Creating backup LSP, link down</b>	A <b>link down</b> event occurred, and traffic is being switched over to the bypass LSP.	<b>extensive</b>
<b>Deleting backup LSP, protected LSP restored</b>	Link has come back up and the LSP has been restored. Because the backup LSP is no longer needed, it is deleted.	<b>extensive</b>
<b>Path mtu</b>	Displays the value of the path MTU received from the network (through signaling) and the value used for forwarding. This value is only displayed on ingress routing devices with the <b>allow-fragmentation</b> statement configured at the <b>[edit protocols mpls path-mtu]</b> hierarchy level. If there is a detour LSP, the path MTU for the detour is also displayed.	<b>detail</b>

Table 67: show rsvp session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Egress protection PLR as protector	RSVP state on the Protector or the point-of-local-repair (PLR) routing device: <ul style="list-style-type: none"> <li><b>Active</b>— Egress protection is available at the Protector or the PLR routing device.</li> <li><b>In Use</b>— Local repair has been completed.</li> </ul>	detail
PATH rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor.	detail
Adspec	MTU signaled from the ingress routing device to the egress routing device by means of the adspec object.	detail
PATH sentto	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor (or peer-name in the GMPLS LSP case), and number of packets sent to the downstream routing device.	detail
Explct route	Explicit route for the session. Normally this value will be the same as that of record route. Differences indicate that path rerouting has occurred, typically during fast reroute.	detail
Record route	Recorded route for the session, taken from the record route object. Normally this value will be the same as that of explct route. Differences indicate that path rerouting has occurred, typically during fast reroute.	detail

## Sample Output

### show rsvp session

```

user@host> show rsvp session
Ingress RSVP: 1 sessions
To          From          State  Rt  Style Labelin Labelout LSPName
10.255.245.214 10.255.245.212 AdminDn 0 1 FF - 22293 LSP Bidir
Total 1 displayed, Up 1, Down 0

Egress RSVP: 2 sessions
To          From          State  Rt  Style Labelin Labelout LSPName
10.255.245.194 10.255.245.195 Up    0 1 FF 39811 - Gpro3-ba Bidir
10.255.245.194 10.255.245.195 Up    0 1 FF 3 - pro3-ba
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To          From          State  Rt  Style Labelin Labelout LSPName
10.255.245.198 10.255.245.197 Up    0 1 SE 100000 3 pro3-de
Total 1 displayed, Up 1, Down 0

```

### show rsvp session statistics

```

user@host> show rsvp session statistics
Ingress RSVP: 2 sessions
To          From          State  Packets  Bytes  LSPName
10.255.245.24 10.255.245.22 Up    0        0  pro3-bd
10.255.245.24 10.255.245.22 Up   44868 2333136 pro3-bd-2

```

```

Total 2 displayed, Up 2, Down 0
Egress RSVP: 2 sessions
To          From          State   Packets   Bytes   LSPName
10.255.245.22 10.255.245.24   Up      0         0    pro3-db
10.255.245.22 10.255.245.24   Up      0         0    pro3-db-2
Total 2 displayed, Up 2, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

```

#### show rsvp session detail

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPName: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left: -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt

```

#### show rsvp session detail (When Egress Protection is in Standby Mode)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPName: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left: -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  Egress protection PLR as protector: Active
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt

```

#### show rsvp session detail (When Egress Protection is in Effect During a Local Repair)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Down, ActiveRoute: 0
  LSPName: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left: -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  Egress protection PLR as protector: In Use
  PATH rcvfrom: localclient

```

```

Adspec: sent MTU 1500
PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt

```

#### show rsvp session detail (Path MTU Output Field)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
10.255.245.3
  From: 10.255.245.5, LSPstate: Up, ActiveRoute: 3
  LSPname: to-c, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100432
  Resv style: 1 FF, Label in: -, Label out: 100432
  Time left: -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
  Port number: sender 1 receiver 57843 protocol 0
  FastReroute desired
  PATH rcvfrom: localclient
  Adspec: sent MTU 4470
  Path mtu: received 4470, using 4458 for forwarding
  PATH sentto: 192.168.37.89 (so-0/2/3.0) 11 pkts
  RESV rcvfrom: 192.168.37.89 (so-0/2/3.0) 10 pkts
  Explct route: 192.168.37.89
  Record route: <self> 192.168.37.89 192.168.37.87
    Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
    Detour adspec: sent MTU 1512
    Path mtu: received 1512, using 1500 for forwarding

```

#### show rsvp session detail (GMPLS)

```

user@host> show rsvp session detail
Ingress RSVP: 1 sessions
192.168.4.1
  From: 192.168.1.1, LSPstate: Dn, ActiveRoute: 0
  LSPname: gmpls-r1-to-r3, LSPpath: Primary
  Bidirectional, Upstream label in: 21253, Upstream label out: -
  Suggested label received: -, Suggested label sent: 21253
  Recovery label received: -, Recovery label sent: -
  Resv style: 0 -, Label in: -, Label out: -
  Time left: -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
  Port number: sender 2 receiver 46115 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH MTU: received 0
  PATH sentto: 10.35.1.5 (so-0/2/3.0) 11 pkts
  Explct route: 100.100.100.100 93.93.93.93
  Record route: <self> 100.100.100.100 93.93.93.93
  Total 1 displayed, Up 0, Down 1
  Egress RSVP: 0 sessions
  Total 0 displayed, Up 0, Down 0
  Transit RSVP: 0 sessions
  Total 0 displayed, Up 0, Down 0

```

#### show rsvp session extensive

```

user@host> show rsvp session extensive
Ingress RSVP: 1 sessions

192.168.0.4
  From: 192.168.0.5, LSPstate: Up, ActiveRoute: 0

```

```

LSPname: E-D, LSPpath: Primary
LSPtype: Static Configured
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: 299808
Resv style: 1 FF, Label in: -, Label out: 299808
Time left: -, Since: Thu Sep 20 15:54:20 2012
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 2 receiver 61576 protocol 0
Attrib flags: Non-PHP
PATH rcvfrom: localclient
Adspec: sent MTU 1500
Path MTU: received 1500
PATH sentto: 10.0.0.18 (lt-1/2/0.17) 41 pkts
RESV rcvfrom: 10.0.0.18 (lt-1/2/0.17) 40 pkts
Explct route: 10.0.0.18 10.0.0.22
Record route: <self> 10.0.0.18 10.0.0.22
Total 1 displayed, Up 1, Down 0

```

Egress RSVP: 1 sessions

192.168.0.5

```

From: 192.168.0.4, LSPstate: Up, ActiveRoute: 0
LSPname: E-D, LSPpath: Primary
Suggested label received: -, Suggested label sent: -
Recovery label received: -, Recovery label sent: -
Resv style: 1 FF, Label in: 3, Label out: -
Time left: 140, Since: Thu Sep 20 15:52:10 2012
Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
Port number: sender 1 receiver 49601 protocol 0
PATH rcvfrom: 10.0.0.18 (lt-1/2/0.17) 44 pkts
Adspec: received MTU 1500
PATH sentto: localclient
RESV rcvfrom: localclient
Record route: 10.0.0.22 10.0.0.18 <self>
Total 1 displayed, Up 1, Down 0

```

Transit RSVP: 0 sessions

Total 0 displayed, Up 0, Down 0

#### show rsvp session p2mp (Ingress Router)

```

user@host> show rsvp session p2mp
Ingress RSVP: 3 sessions
P2MP name: test, P2MP branch count: 1
To      From      State  Rt Style Labelin Labelout LSPname
10.255.10.95 10.255.10.2  Up    0  1 SE  -        3 to-pe1
P2MP name: test2, P2MP branch count: 2
To      From      State  Rt Style Labelin Labelout LSPname
10.255.10.23 10.255.10.2  Up    0  1 SE  -        299776 to-pe3
10.255.10.16 10.255.10.2  Up    0  1 SE  -        299776 to-pe4
Total 3 displayed, Up 3, Down 0

```

Egress RSVP: 0 sessions

Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions

Total 0 displayed, Up 0, Down 0

#### show rsvp session p2mp (Transit Router)

```

user@host> show rsvp session p2mp

```

## Ingress RSVP: 1 sessions

P2MP name: test, P2MP branch count: 1

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.255.10.23	10.255.10.95	Up	0	1 SE	-	299792	to-pe2

Total 1 displayed, Up 1, Down 0

## Egress RSVP: 1 sessions

P2MP name: test, P2MP branch count: 1

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.255.10.95	10.255.10.2	Up	0	1 SE	3	-	to-pe1

Total 1 displayed, Up 1, Down 0

## Transit RSVP: 2 sessions

P2MP name: test2, P2MP branch count: 2

To	From	State	Rt	Style	Labelin	Labelout	LSPname
10.255.10.23	10.255.10.2	Up	0	1 SE	299776	299808	to-pe3
10.255.10.16	10.255.10.2	Up	0	1 SE	299776	299856	to-pe4

Total 2 displayed, Up 2, Down 0



## show rsvp statistics

<b>List of Syntax</b>	<a href="#">Syntax on page 585</a> <a href="#">Syntax (EX Series Switches) on page 585</a>
<b>Syntax</b>	<pre>show rsvp statistics &lt;instance <i>instance-name</i>&gt; &lt;logical-system (all   <i>logical-system-name</i>)&gt;</pre>
<b>Syntax (EX Series Switches)</b>	show rsvp statistics
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p><b>instance <i>instance-name</i></b> option added in Junos OS Release 15.1.</p>
<b>Description</b>	Display Resource Reservation Protocol (RSVP) packet and error statistics.
<b>Options</b>	<p><b>none</b>—Display RSVP packet and error statistics.</p> <p><b>instance <i>instance-name</i></b>—(Optional) Display RSVP packet and error statistics for the specified instance. If <b><i>instance-name</i></b> is omitted, RSVP statistics is displayed for the master instance.</p> <p><b>logical-system (all   <i>logical-system-name</i>)</b>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear rsvp statistics on page 556</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show rsvp statistics on page 588</a>
<b>Output Fields</b>	<p><a href="#">Table 68 on page 585</a> describes the output fields for the <b>show rsvp statistics</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 68: show rsvp statistics Output Fields**

Field Name	Field Description
<b>Packet Type</b>	Statistics about different RSVP messages.
<b>Total Sent</b>	Total number of packets sent since RSVP was enabled.
<b>Total Received</b>	Total number of packets received since RSVP was enabled.
<b>Last 5 seconds Sent</b>	Total number of packets sent in the last 5 seconds.
<b>Last 5 seconds Received</b>	Number of packets received in the last 5 seconds.

Table 68: show rsvp statistics Output Fields (*continued*)

Field Name	Field Description
<b>Path</b>	Statistics about Path messages, which are sent from the RSVP sender along the data paths and which store path state information in each node along the path.
<b>PathErr</b>	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.
<b>PathTear</b>	Statistics about PathTear messages, which remove path states and dependent reservation states in any routing devices along a path.
<b>Resv FF</b>	Statistics about fixed-filter reservation style messages, which consist of distinct reservations among explicit senders.
<b>Resv WF</b>	Statistics about wildcard-filter reservation style messages, which consist of shared reservations among wildcard senders.
<b>Res SE</b>	Statistics about shared-explicit reservation style messages, which consist of shared reservations among explicit senders.
<b>ResvErr</b>	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.
<b>ResvTear</b>	Statistics about ResvTear messages, which remove reservation states along a path.
<b>ResvConf</b>	Statistics about ResvConfirm messages, which are responses to confirm a reservation request.
<b>Ack</b>	Acknowledge message for refresh reductions.
<b>SRefresh</b>	Summary refresh messages.
<b>Hello</b>	Number of RSVP hello packets that have been sent to and received from the neighbor.
<b>EndtoEnd RSVP</b>	Statistics for the number of End-to-end RSVP messages.
<b>Errors</b>	Statistics about errored RSVP packets.
<b>Rcv pkt bad length</b>	The packet was not processed because its length is inappropriate.
<b>Rcv pkt unknown type</b>	The packet is not one of the well-known RSVP types, as defined in RFC 2205, <i>Resource ReSerVation Protocol (RSVP)</i> .
<b>Rcv pkt bad version</b>	The packet is not an RSVP version 1 packet.
<b>Rcv pkt auth fail</b>	The packet failed authentication checks.
<b>Rcv pkt bad checksum</b>	The RSVP checksum check failed.
<b>Rcv pkt bad format</b>	General packet processing failed because the packet was badly formed.
<b>Memory allocation fail</b>	An internal resource failure occurred.

Table 68: show rsvp statistics Output Fields (*continued*)

Field Name	Field Description
No path information	A reservation was received, but no sender is active.
Resv style conflict	The same session contains inconsistent reservation styles.
Port conflict	There were inconsistent port numbers for the same session.
Resv no interface	An interface for the receive reservation packets cannot be located.
PathErr to client	Number of PathErr packets delivered to the local client.
ResvErr to client	Number of ResvErr packets delivered to the local client.
Path timeout	Number of times the sender timed out because the path was removed.
Resv timeout	Number of times the receiver timed out because the reservation was removed.
Message out-of-order	Records the number of RSVP incoming messages that are considered out of order. This is detected from the message ID object's sequence number.
Unknown ack msg	A neighboring routing device replies with an ACK object that contains an unknown message ID. This can indicate a message ID handshake problem. For example, a router receives an ACK for message IDs 1, 2, and 3. However, it only has state for message IDs 1 and 3. The router increments the unknown ack counter by 1.
Recv nack	If a neighboring router receives an unknown message ID in an RSVP refresh message, the router sends a Resv nack message back to the sender. This can happen if that neighbor has been rebooted. For this case, the router sends a regular RSVP refresh message to recover the state and start the message-ID handshake process again.
Recv duplicated msg-id	Number of times the same message ID is used by two different RSVP messages. This duplication is usually caused when a neighboring routing device restarts.
No TE-link to rcv Hop	Counter of packets discarded because a TE link was not found.
Rcv pkt disabled interface	Number of RSVP packets received on an interface that is not enabled for RSVP.
Transmit buffer full	Number of times the buffer for assembling an outgoing RSVP message was not large enough.
Transmit failure	Number of times the RSVP task failed to send out a packet.
Receive failure	Number of times the RSVP task failed to read an incoming packet.
P2MP RESV discarded by appl	Number of Resv messages discarded because the MPLS label is not valid for the P2MP LSP application.
Rate limit	Number of RSVP packets dropped due to rate limiting.

Table 68: show rsvp statistics Output Fields (*continued*)

Field Name	Field Description
<b>Err msg loop detected</b>	Number of RSVP error messages that have looped back to their originator. This is detected by checking the error node address in the ERROR_SPEC object.

## Sample Output

### show rsvp statistics

```

user@host> show rsvp statistics
  PacketType          Sent      Received      Last 5 seconds
                                     Sent      Received
  Path                355         408           0           0
  PathErr              2          13           0           0
  PathTear            101        139           0           0
  Resv FF              0           0           0           0
  Resv WF              0           0           0           0
  Resv SE             419        225           0           0
  ResvErr              0           0           0           0
  ResvTear             0          13           0           0
  ResvConf             0           0           0           0
  Ack                  682       1414           0           0
  SRefresh            395198     236030         5           2
  Hello               578809     578221         4           4
  EndtoEnd RSVP        0           0           0           0

  Errors              Total      Last 5 seconds
  Rcv pkt bad length      0           0
  Rcv pkt unknown type    0           0
  Rcv pkt bad version     0           0
  Rcv pkt auth fail       0           0
  Rcv pkt bad checksum    0           0
  Rcv pkt bad format      0           0
  Memory allocation fail  0           0
  No path information     10          0
  Resv style conflict     0           0
  Port conflict           0           0
  Resv no interface       0           0
  PathErr to client       38          0
  ResvErr to client       0           0
  Path timeout            8           0
  Resv timeout            57          0
  Message out-of-order    0           0
  Unknown ack msg         2978        0
  Recv nack               86          0
  Recv duplicated msg-id   5           0
  No TE-link to rcv Hop   0           0
  Rcv pkt disabled interface 0           0
  Transmit buffer full    0           0
  Transmit failure        0           0
  Receive failure         0           0
  P2MP RESV discarded by appl 0           0
  Rate limit              306          0
  Err msg loop detected    0           0

```

## show rsvp version

<b>List of Syntax</b>	<a href="#">Syntax on page 589</a> <a href="#">Syntax (EX Series Switches) on page 589</a>
<b>Syntax</b>	show rsvp version <logical-system (all   <i>logical-system-name</i> )>
<b>Syntax (EX Series Switches)</b>	show rsvp version
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
<b>Description</b>	Display information about the Resource Reservation Protocol (RSVP) protocol settings, such as the version of the RSVP software, the refresh timer and keep multiplier, and local RSVP graceful restart capabilities on a routing device.
<b>Options</b>	<b>none</b> —Display RSVP protocol settings.  <b>logical-system (all   <i>logical-system-name</i>)</b> —(Optional) Perform this operation on all logical systems or on a particular logical system.
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show rsvp version on page 590</a>
<b>Output Fields</b>	<a href="#">Table 69 on page 589</a> describes the output fields for the <b>show rsvp version</b> command. Output fields are listed in the approximate order in which they appear.

**Table 69: show rsvp version Output Fields**

Field Name	Field Description
Resource ReSerVation Protocol, version	RSVP software version.
RSVP protocol	Status of RSVP: <b>Enabled</b> or <b>Disabled</b> .
R(refresh timer)	Configured time interval used to generate periodic RSVP messages.
K(keep multiplier)	Number of RSVP messages that can be lost before an RSVP state is declared stale.
Preemption	Currently configured preemption capability: <b>Aggressive</b> , <b>Disabled</b> , or <b>Normal</b> . The default is <b>Normal</b> .
Soft-preemption cleanup	Time, in seconds, that an LSP is kept after it has been soft preempted. This is a global property of the RSVP protocol.
Graceful deleting timeout	Currently configured value for the <b>graceful-deletion-timeout</b> statement. The router that initiates the graceful deletion procedure for an RSVP session waits for the graceful deletion timeout interval to ensure that all routers along the path (especially the ingress and egress routers) have prepared for the LSP to be taken down.

Table 69: show rsvp version Output Fields (*continued*)

Field Name	Field Description
<b>NSR Mode</b>	Status of the nonstop active routing feature for RSVP on the restarting device: <b>Disabled</b> , <b>Enabled/Master</b> , or <b>Enabled/Standby</b> .
<b>NSR State</b>	<p>State of the nonstop active routing feature for RSVP on the restarting device.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> <li>• <b>Idle</b></li> <li>• <b>TE-link sync complete</b></li> <li>• <b>Neighbor sync complete</b></li> <li>• <b>Path state sync complete</b></li> <li>• <b>Resv state sync complete</b></li> <li>• <b>Bypass sync complete</b></li> <li>• <b>Init sync complete</b></li> </ul>
<b>Setup protection</b>	Status of point-to-point and point-to-multipoint LSP setup protection configuration on the device: <b>Enabled</b> or <b>Disabled</b>
<b>Graceful restart</b>	Status of the graceful restart feature for RSVP on the restarting routing device: <b>Enabled</b> or <b>Disabled</b> .
<b>Restart helper mode</b>	Status of the helper mode feature: <b>Enabled</b> or <b>Disabled</b> . When this feature is enabled, the restarting routing device can help the neighbor with its RSVP restart procedures.
<b>Maximum helper restart time</b>	Number of milliseconds (ms) configured for the maximum helper restart time. The maximum helper restart time is the length of time the routing device waits before declaring that an RSVP neighbor attempting to restart gracefully is down.
<b>Maximum helper recovery time</b>	Number of milliseconds configured for the maximum helper recovery time. The maximum helper recovery time is the amount of time the routing device maintains the state of an RSVP neighbor attempting to restart gracefully.
<b>Restart time</b>	Number of milliseconds that a neighbor waits to receive a Hello message from the restarting node before declaring the node dead and deleting the states.
<b>Recovery time</b>	Number of milliseconds during which the restarting node attempts to recover its lost states with help from its neighbors. Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.
<b>P2p transit LSP nexthop mode</b>	Point-to-point transit LSP nexthop mode on PTX Series devices. The possible values are <b>Chained</b> or <b>Unchained</b>
<b>P2mp transit LSP nexthop mode</b>	Point-to-multipoint transit LSP nexthop mode on PTX Series devices. The possible values are <b>Chained</b> or <b>Unchained</b>

## Sample Output

### show rsvp version

```
user@host> show rsvp version
```

```
Resource ReSerVation Protocol, version 1. rfc2205
  RSVP protocol:           Enabled
  R(refresh timer):        30 seconds
  K(keep multiplier):      3
  Preemption:              Normal
  Soft-preemption cleanup:  30 seconds
  Graceful deletion timeout: 30 seconds
  NSR mode:                 Enabled/Master
  NSR state:                Init sync complete
  Setup protection:        Disabled
  Graceful restart:        Disabled
  Restart helper mode:     Enabled
  Maximum helper restart time: 20000 msec
  Maximum helper recovery time: 180000 msec
  Restart time:            0 msec
  P2p transit LSP nexthop mode: Unchained
  P2mp transit LSP nexthop mode: Unchained
```

## traceroute mpls rsvp

---

<b>Syntax</b>	<code>traceroute mpls &lt;rsvp&gt; <i>lsp-name</i></code> <code>&lt;detail&gt;</code> <code>&lt;egress&gt;</code> <code>&lt;exp&gt;</code> <code>&lt;logical-system&gt;</code> <code>&lt;multipoint&gt;</code> <code>&lt;no-resolve&gt;</code> <code>&lt;retries&gt;</code> <code>&lt;source <i>source-address</i>&gt;</code> <code>&lt;ttl&gt;</code>
<b>Release Information</b>	Command introduced in Junos OS Release 9.2. <code>egress</code> , <code>multipoint</code> , and <code>ttl</code> options added in Junos OS Release 11.2.
<b>Description</b>	Trace route to a remote host for an MPLS LSP signaled by RSVP. Use <b>traceroute mpls rsvp</b> as a debugging tool to locate MPLS label-switched path (LSP) forwarding issues in a network. (Currently supported for IPv4 packets only.)
<b>Options</b>	<p><b><i>lsp-name</i></b>—Specify the name of the LSP to be traced.</p> <p><b><code>detail</code></b>—(Optional) Display detailed output.</p> <p><b><code>egress</code></b>—(Optional) Request that a specific point-to-multipoint egress node reply to the trace route. The trace route would follow the associated sub-LSP to the egress node.</p> <p><b><code>exp</code></b>—(Optional) Specify the class of service to use when sending probes. The range of values is 0 through 7. The default value is 7.</p> <p><b><code>logical-system</code></b>—(Optional) Specify the name of the logical system for the traceroute attempt.</p> <p><b><code>multipoint</code></b>—(Optional) Perform a trace route on a point-to-multipoint LSP.</p> <p><b><code>no-resolve</code></b>—(Optional) Specify not to resolve the hostname that corresponds to the IP address.</p> <p><b><code>retries</code></b>—(Optional) Specify the number of times to resend probe. The range of values is 1 through 9. The default value is 3.</p> <p><b><code>source <i>source-address</i></code></b>—(Optional) Specify the source address of the outgoing traceroute packets.</p> <p><b><code>ttl</code></b>—(Optional) Specify the number of hops to follow before forcing the trace route to quit.</p>
<b>Required Privilege Level</b>	network
<b>List of Sample Output</b>	<a href="#">traceroute mpls rsvp on page 594</a> <a href="#">traceroute mpls rsvp detail on page 594</a>



[traceroute mpls rsvp multipoint \(branch node for sub-LSPs\) on page 595](#)

[traceroute mpls rsvp multipoint \(single-hop sub-LSPs\) on page 595](#)

**Output Fields** Table 70 on page 593 describes the output fields for the **traceroute mpls rsvp *lsp-name*** and **traceroute mpls rsvp *lsp-name* detail** commands. Output fields are listed in the approximate order in which they appear.

**Table 70: traceroute mpls rsvp Output Fields**

Field Name	Field Description	Level of Output
Probe options	Probe options specified in the <b>traceroute mpls rsvp <i>lsp-name</i></b> command.	all levels
ttl	Time-to-live value of the labeled packet.	none specified
Label	MPLS label used to forward the packets along the LSP.	none specified
Protocol	Signaling protocol used. For this command, it is RSVP-TE.	none specified
Address	Address of the next hop.	none specified
Previous Hop	Address of the previous hop. Previous hop address of the first hop is null.	none specified
Probe status	Forwarding status from the first hop to the last-hop label-switching router (egress point in the label-switched paths). Displays <b>Success</b> if the trace to a hop is successful or <b>Egress</b> if the trace has reached the last router on the path.	none specified
Hop	Address of the hops in the label-switched path from the first hop to the last hop. Depth indicates the level of the hop.	<b>detail</b>
Parent	Address of the previous hop. Parent value for the first hop is null.	<b>detail</b>
Return Code	Return code for reporting the result of processing the echo request by the receiver.	<b>detail</b>
Sender timestamp	Displays the timestamp when the MPLS echo request is sent to the next hop.	<b>detail</b>
Receiver timestamp	Timestamp when the echo request from the previous hop is received and acknowledged with an echo response by the next hop.	<b>detail</b>
Response time	Time for the echo request to reach the receiver.	<b>detail</b>
MTU	Size of the largest packet that includes the label stack forwarded to the next hop.	<b>detail</b>

Table 70: traceroute mpls rsvp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Multipath type	Labels or addresses used by the specified multipath type. If multipaths are not used, the value is none.	<b>detail</b>
Label stack	Label stack used to forward the packet.	<b>detail</b>
Path	Displays the sub-lsp path number for this traceroute, the interface used, and the destination address.	all levels

## Sample Output

### traceroute mpls rsvp

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta
```

```
Probe options: retries 3, exp 7
```

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	299792	RSVP-TE	192.168.1.2	(null)	Success
2	299803	RSVP-TE	192.168.2.3	192.168.1.2	Success
3	3	RSVP-TE	192.168.3.4	192.168.2.3	Egress

```
Path 1 via ge-0/0/0.1 destination 127.0.0.64
```

### traceroute mpls rsvp detail

```
user@host> traceroute mpls rsvp lsp-chicago-atlanta detail
```

```
Probe options: retries 3, exp 7
```

```
Hop 192.168.1.2 Depth 1
```

```
Probe status: Success
```

```
Parent: (null)
```

```
Return code: Label-switched at stack-depth 1
```

```
Sender timestamp: 2008-04-17 09:35:27 EDT 400.88 msec
```

```
Receiver timestamp: 2008-04-17 09:35:27 EDT 427.87 msec
```

```
Response time: 26.99 msec
```

```
MTU: Unknown
```

```
Multipath type: IP bitmask
```

```
Address Range 1: 127.0.0.64 ~ 127.0.0.127
```

```
Label Stack:
```

```
Label 1 Value 299792 Protocol RSVP-TE
```

```
Hop 192.168.2.3 Depth 2
```

```
Probe status: Success
```

```
Parent: 192.168.1.2
```

```
Return code: Upstream interface index unknown label-switched at stack-depth
```

```
1
```

```
Sender timestamp: 2008-04-17 09:35:27 EDT 522.13 msec
```

```
Receiver timestamp: 2008-04-17 09:35:27 EDT 548.69 msec
```

```
Response time: 26.55 msec
```

```
MTU: 1518
```

```
Multipath type: IP bitmask
```

```
Address Range 1: 127.0.0.64 ~ 127.0.0.127
```

```
Label Stack:
```

```
Label 1 Value 299803 Protocol RSVP-TE
```

**traceroute mpls rsvp multipoint (branch node for sub-LSPs)**

The following traceroute output is for a point-to-multipoint LSP where the penultimate node is a branch node for the sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
Probe options: retries 3, exp 7
```

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	300000	RSVP-TE	81.1.2.2	(null)	Success
2	299968	RSVP-TE	81.2.3.3	81.1.2.2	Success
3	299952	RSVP-TE	81.3.4.4	81.2.3.3	Success
4	299920	RSVP-TE	81.4.6.6	81.3.4.4	Egress

Path 1 via lt-1/2/0.102 destination 127.0.0.64

ttl	Label	Protocol	Address	Previous Hop	Probe Status
4	299920	RSVP-TE	81.4.5.5	81.3.4.4	Egress

Path 2 via lt-1/2/0.102 destination 127.0.0.64

**traceroute mpls rsvp multipoint (single-hop sub-LSPs)**

The following traceroute output is for a point-to-multipoint LSP with multiple single-hop sub-LSPs.

```
user@host> traceroute mpls rsvp multipoint p2mplsp
Probe options: retries 3, exp 7
```

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	0	RSVP-TE	81.1.2.2	(null)	Egress

Path 1 via lt-1/2/0.102 destination 127.0.0.64

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	0	RSVP-TE	81.1.8.8	(null)	Egress

Path 2 via lt-1/2/0.108 destination 127.0.0.64

ttl	Label	Protocol	Address	Previous Hop	Probe Status
1	0	RSVP-TE	81.1.9.9	(null)	Egress

Path 3 via lt-1/2/0.109 destination 127.0.0.64

